arXiv:1510.06813v2 [q-fin.EC] 22 Jun 2016

Analysis of Markovian Competitive Situations using Nonatomic Games—the Shock-driven Case and its Dynamic Pricing Application

Jian Yang

Department of Management Science and Information Systems

Business School, Rutgers University

Newark, NJ 07102

Email: jyang@business.rutgers.edu

June 2016

Abstract

We show that equilibria derived for nonatomic games (NGs) can be used by large finite games to achieve near-equilibrium performances. Rather than single-period games, we deal with dynamic games in which the evolution of a player's state is influenced by his own action as well as other players' states and actions. We focus on the case where both random state transitions and random actions are driven by independently generated shocks. The NG equilibria we consider are random state-to-action maps. The simple NG equilibria are adoptable to a variety of real situations where awareness of other players' states can be anywhere between full and none. Transient results here also form the basis for a link between an NG's stationary equilibrium (SE) and good stationary profiles for large finite games. Our approach works well for certain dynamic pricing games, both without and with production.

Keywords: Nonatomic Game; Markov Equilibrium; Dynamic Pricing

1 Introduction

Many multi-period competitive situations, as first noted by Shapley [31], involve randomly-evolving player states that affect players' payoffs. When making a decision, a player has to be concerned with not only what states other players are in and how other players will act, but also how his and others' states and actions will influence the future evolution of the players' states. Another complicating factor is that players may have zero, partial, or full knowledge of other players' states before they take their actions in each period. The task of analyzing these dynamic games is certainly challenging. As an example of such competitive situations, consider a dynamic pricing game. Multiple firms start a fixed time horizon with stocks of the same product. Each of them is bent on using pricing to influence demand and earn the highest revenue from selling their respective stocks. In any given period, a firm is aware of its own current inventory level but not the levels of others. Yet, demand arrival to the firm is both random and influenced by not only its own price, but also prices charged by other firms. In spite of its practical relevance, this dynamic pricing game has not received satisfactory treatment due probably to its obvious complexity.

The ultimate goal with such a Markovian game lies in identifying an equilibrium action plan that will earn each player the highest total payoff when other players adhere to the plan. But even in the stationary setting, known equilibria come in quite complicated forms that for real implementation, demand a high degree of coordination among players; see, e.g., Mertens and Parthasrathy [25], Duffie, Geanakoplos, Mas-Colell, and McLennan [8], and Solan [32]. An exact equilibrium thus being too lofty to attain, we propose that equilibria be reached asymptotically as the number of players grows, on the premise that the game's nonatomic-game (NG) counterpart is analyzable. In the latter, a continuum of players are in competition, none of whom having any discernible influence on any other players and yet all players in aggregation hold sway on players' payoffs and state evolutions. The key advantage of such a game is that its state distribution will evolve in a deterministic fashion. This results in the relatively simple form taken by the NG's equilibria x: the pure or mixed action plan $x_t(s_t)$, though dependent on the time period t and his own individual state s_t, is insensitive to whatever portion of the overall state distribution that the player can observe.

When an NG equilibrium is handy, we show that it can be used on the original finite Markovian game to serve our intended purpose. Relying on intermediate results stemming from the weak Law of Large Numbers (LLN) concerning empirical distributions, we establish two main results. In Theorem 1, we show that the empirical distribution of players' states,

which is itself random in the finite game, will nevertheless converge in probability to the deterministic distribution as predicted for the NG counterpart when the number of players grows to infinity. This convergence paves way for Theorem 2, which states that players can use the observation-blind NG equilibrium in the finite-player situation and gain an average performance that is ever harder to beat as the number of players grows. In both results, the "average" on players' states is assessed on the state distribution prevailing in either the NG or the finite game. After assuming time-invariant payoff and transition functions, as well as fixed discountings over time and an infinite time horizon, we obtain a stationary setting. For this, we establish Theorem 3, effectively our affirmative answer to whether stationary equilibria (SE) studied in past literature can be useful to large finite games.

The above theory will be most useful when the NG counterpart is relatively easy to deal with in comparison to the corresponding finite games. Besides evidence in literature, this point is further buttressed by the dynamic pricing game mentioned earlier. Our analysis of it demonstrates the usefulness of the transient result Theorem 2. We extend the game by considering locked-in production, wherein every firm uses production to bring its inventory back up to a pre-determined level whenever it becomes empty. This game is asymptotically analyzable thanks to the stationary result Theorem 3.

As our foremost contribution, we established one more link between NGs and their finite-game counterparts. Previously, links were mostly established for single-period games, special multi-period games without individual states, or games exhibiting stationary features. The introduction of information-carrying individual states allow in for proper treatment a much wider swath of applicable situations involving present-future tradeoffs and transient properties. Comparing to the earlier work Yang [38] which dealt with the NG-finite link for Markovian games as well, the current paper treats more general, non-discrete state and action spaces. As a tradeoff, we are compelled to assume that random shocks drive both decision making and state evolution. Our analysis later, however, show that these do not much restrict the generality of our results. Besides, methodologies employed are vastly different. The earlier paper faced sampling from joint probabilities that are of the non-product type, which placed limitations on state and action spaces that could be handled. Thanks to the presence of shocks in the current work, we can avoid sampling from joint probabilities of the non-product type and resort to tools pertinent to the tightness of probabilities which only require certain spaces to be complete metric.

Secondly, we demonstrated that the usefulness of SEs to large finite games stems from

more fundamental properties possessed by transient NG equilibria. Last but not least, we added to the revenue management literature a new way to tackle the notoriously difficult dynamic pricing game in which firms are capable of reacting in real time to their changing inventory levels and are yet blind to other firms' inventory levels. The shocks involved in decision making allow the randomization of firms' inventory-to-price maps. This realistic feature makes it easy for firms to reach equilibirum in the NG regime. Indeed, our dynamic pricing game only requires that the demand-price relationship be continuous in the joint type-price distribution of other firms. No concavity- or monotonicity-related assumption which could reduce the appeal of price-competition results is needed here.

Here is our plan for the remainder of the paper. We spend Section 2 on a survey of related research and Section 3 on basic model primitives. The nonatomic game is introduced in Section 4, while finite games are treated in Section 5. We present the main transient convergence results in Section 6. These are used in Section 7 to establish a link between SEs and large finite games with stationary features. We apply our results to dynamic price competition among heterogeneous firms in Section 8. Further discussion is made in Section 9, while the paper is concluded in Section 10.

2 Literature Survey

NGs are often easier to analyze than their finite counterparts, because in them, the action of an individual player has no impact on payoffs and future state evolutions of the other players. Therefore, they are often used as proxies of real competitive systems in economic studies; see, e.g., Aumann [5] and Reny and Perry [28]. Systematic research on NG started with Schmeidler [30]. He formulated a single-period semi-anonymous NG, wherein the joint distribution of other players' types and actions may affect any given player's payoff. When the action space is finite, Schmeidler established the existence of pure equilibria when the game becomes anonymous, so that only the distribution of other players' actions matters. Mas-Colell [24] showed the existence of distributional equilibria in anonymous NGs with compact action spaces. Khan, Rath, and Sun [21] identified a certain limit to which Schmeidler's result can be extended. Links between NGs and their finite counterparts were covered in Green [13], Housman [15], Carmona [7], Kalai [20], Al-Najjar [4], and Yang [37].

This paper differs from the above by its focus on multi-period games. For such games without individual states that allow past actions to impact future gains, Green [12], Sabourian

[29], and Al-Najjar and Smorodinsky [3] showed that equilibria for large games are nearly myopic. With individual states that inherit traces of past actions, the games we study pose new challenges. An NG equilibrium for our situation is certainly not myopic as it takes into account the current action's future consequences. Rather, it is insensitive to real-time observations made on other players' states. We succeed in showing that such a simple action plan can be used in finite situations with randomly evolving state distributions of which a player may have zero, partial, or full knowledge, and still generate decent payoffs. The type of NGs we deal with are similar to sequential anonymous games studied by Jovanovic and Rosenthal [18], who established existence of distributional equilibria. The result was generalized to games involving aggregate shocks by Bergin and Bernhardt [6]. Different from these papers, we work on the link between NGs and finite games, not the NGs themselves.

In their effort to simplify dynamic games, some authors went further than silencing individual player influences as done through the NG approach. In addition, they strove for the so-called stationary equilibria (SE), that further stressed the long-run steady-state nature of individual action plans and system states; see, e.g., Hopenhayn [14] and Adlakha and Johari [1]. The oblivious equilibrium (OE) concept as proposed by Weintraub, Benkard, and van Roy [34], though accounting for impacts of large players, took the same stationary approach by letting firms beware of only long-run average system states. We caution that the implicit stationarity of SE or OE renders it inappropriate for applications that are transient by nature; for instance, the dynamic pricing game to be studied wherein the inventory level of every firm can only decrease over time.

Some recent works also contributed to the understandings of links between equilibria of infinite-player games and their finite-player brethren. Weintraub, Benkard, and van Roy [35] did so for a setting where long-run average system state can be defined. Adlakha, Johari, and Weintraub [2] established the existence of SE and achieved a similar conclusion by using only exogenous conditions on model primitives. Weintraub et al. [33] studied nonstationary oblivious equilibria (NOE) that capture transient behaviors of players, and showed their usefulness in finite-player situations by relying on a "light-tail" condition on players' state distributions similar to that used in [35].

Huang, Malhame, and Caines [16] dealt with a continuous-time multi-player system where independent diffusion processes provide random drivers. They reached equilibria in the nonatomic limit, and derived asymptotic results as the number of players becomes large. In the work, other players impact a given player through linear functionals of the state

distribution they form; meanwhile, their actions play no direct role. Our discrete-time framework afforded us almost full generality when it comes to other players' impacts on any given player's payoffs and state transitions—it is the joint state-action distribution that forms the environment faced by an individual player. As already mentioned, while this paper tackles the case where exogenous shocks drive state evolution and decision making, Yang [38] dealt with the setting where such shocks are not necessarily identifiable; however, technical challenges faced there forced state and action spaces to be discrete.

The dynamic pricing example has its own practical significance. In today's marketplace, companies are often obliged to take the competition element into their revenue management considerations. The permeation of Internet access among virtually all homes and businesses brings price competition to an even more intense, global level. Sites that pool dynamically changing prices from tens or even hundreds of providers and let customers choose the most suitable ones for themselves include Expedia, Hotwire, Orbitz, and Travelocity, just to name a few. The competition is heated in several industries and is most manifest in the hospitality industry and the second market of ticket selling.

Yet, research on competitive dynamic pricing has only been sporadic. Perakis and Sood [27] studied such a problem in which each firm reacts most conservatively to the inventory-independent price schedules of its competitors. Xu and Hopp [36] considered a multi-seller dynamic pricing game, where demand arrival is governed by a geometric Brownian motion and only the lowest-priced seller sells. Lin and Sibdari [23] and Levin, McGill, and Nediak [22] resolved equilibrium existence issues for specific dynamic pricing situations; neither revealed structural properties of their equilibria. Gallego and Hu [11] modeled deterministic transient price competition as differential games and applied heavy-traffic approximation to the corresponding stochastic situations. Regarding precisions of the respective approximations, we note that the just-mentioned work relied on there being sufficiently high supply and demand, while we count on there being enough number of competing firms. The earlier work also assumed that opponents' inventories are visible, a point rendered irrelevant by our NG approach. Yang and Xia [39] provided computational support to the current approach. They tackled a continuous-time dynamic pricing game, and found that an NG equilibrium would fare well in an actual competitive situation with as few as 30 firms.

3 Model Primitives

Given a metric space A, we use d_A to denote its metric, $\mathcal{B}(A)$ its Borel σ-field, and $\mathcal{P}(A)$ the set of all probability measures (distributions) on the measurable space $(A, \mathcal{B}(A))$. The space $\mathcal{P}(A)$ is metrized by the Prohorov metric ρ_A, which induces on it the weak topology. Given metric spaces A and B, we use $\mathcal{M}(A, B)$ to represent all measurable functions from A to B. We study multi-period games in which players possess individual states; also, the payoff and random state transition of any given player are influenced by other players' states along with their actions. Exogenously generated random shocks drive both players' decision making processes and states' inter-period transitions.

In particular, we use complete separable metric space S for individual states and separable metric space X for player actions. Every time players are engaged, each is aware of his own individual state $s \in S$. Other players' states and actions are felt in a semi-anonymous fashion: what matters is the joint distribution say $\mu \in \mathcal{P}(S \times X)$ of other players' states and actions. Players will choose their actions according to shocks drawn from a complete separable metric space G. Let the fixed distribution of action-inducing shocks be $\bar{\gamma} \in \mathcal{P}(G)$. After taking their actions, players will each experience another shock from a complete separable metric space I which will help push the evolution of their states. Let the fixed distribution of transition-inducing shocks be $\bar{\iota} \in \mathcal{P}(I)$. The completeness requirements on S, G, and I stem from the need to invoke Lemma 3 in Appendix A. These are not stringent, as multi-dimensional Euclidean spaces, as well as finite and countable sets endowed with the discrete metric, are all complete separable metric spaces.

We consider $\bar{t} = 1, 2, ...$ periods. For $t = 1, ..., \bar{t}$, suppose a player's state $s \in S$, his action $x \in X$, and the joint state-action distribution $\mu \in \mathcal{P}(S \times X)$ he faces, together determine his payoff in period t. Thus, we introduce a bounded payoff function

$$\tilde{f}_t : S \times X \times \mathcal{P}(S \times X) \longrightarrow [-\bar{f}_t, \bar{f}_t], \tag{1}$$

where $\bar{f}_t$ is some positive constant. It is required that $\tilde{f}_t(\cdot, \cdot, \mu) \in \mathcal{M}(S \times X, [-\bar{f}_t, \bar{f}_t])$ for every given distribution $\mu \in \mathcal{P}(S \times X)$. To express the phenomenon that the same player will enter a new state under shock $i \in I$, we introduce mapping

$$\tilde{s}_t : S \times X \times \mathcal{P}(S \times X) \times I \longrightarrow S. \tag{2}$$

We require that $\tilde{s}_t(\cdot, \cdot, \mu, \cdot) \in \mathcal{M}(S \times X \times I, S)$ at every distribution $\mu \in \mathcal{P}(X \times S)$. In Yang [38], state evolution is governed by some $\tilde{g}_t : S \times X \times \mathcal{P}(S \times X) \longrightarrow \mathcal{P}(S)$, so that

$\tilde{g}_t(s, x, \mu|S')$ is the chance that a player's state will come from S' in the next period when he is currently at state s, takes action x, and faces joint state-action distribution μ. Our present case can be understood as one with

$$\tilde{g}_t(s, x, \mu) = \bar{\iota} \cdot (\tilde{s}_t(s, x, \mu, \cdot))^{-1}, \tag{3}$$

meaning that, for any $S' \in \mathcal{B}(S)$,

$$\tilde{g}_t(s, x, \mu|S') = \bar{\iota}(\{i \in I|\tilde{s}_t(s, x, \mu, i) \in S'\}) = \int_I \mathbf{1}(\tilde{s}_t(s, x, \mu, i) \in S') \cdot \bar{\iota}(di), \tag{4}$$

where $\mathbf{1}(\cdot)$ is the indicator function. However, when S is a closed subset of the real line $\Re$, the converse will be true. Then, it would suffice to study the current case with $I = [0, 1]$, $\bar{\iota}$ being uniform, and

$$\tilde{s}_t(s, x, \mu, i) = \inf \{s \in S|\tilde{g}_t(s, x, \mu|(-\infty, s] \cap S) \geq i\}. \tag{5}$$

The action plans we consider are of the form $x_t \in \mathcal{M}(S \times G, X)$. That is, a player will use action $x_t(s, g)$ in period t when he starts with state $s \in S$ and receives action-inducing shock g. This can again be cast through a random action plan $\chi_t : S \longrightarrow \mathcal{P}(X)$ studied in Yang [38], because of the association

$$\chi_t(s) = \bar{\gamma} \cdot (x_t(s, \cdot))^{-1}. \tag{6}$$

This means that, for any $X' \in \mathcal{B}(X)$,

$$\chi_t(s|X') = \bar{\gamma}(\{g \in G|x_t(s, g) \in X'\}) = \int_G \mathbf{1}(x_t(s, g) \in X') \cdot \bar{\gamma}(dg). \tag{7}$$

However, when X is a closed subset of the real line, we can let G be $[0, 1]$, $\bar{\gamma}$ be uniform, and

$$x_t(s, g) = \inf \{x \in X|\chi_t(s|(-\infty, x] \cap X) \geq g\}. \tag{8}$$

Also, we can always use a singleton G and hence degenerate $\bar{\gamma}$ to handle pure action plans and use a singleton I and hence degenerate $\bar{\iota}$ to handle deterministic state transitions.

Let us call any state distribution $\sigma \in \mathcal{P}(S)$ a pre-action environment, because the one formed by other players is what a player sees at the beginning of any period. Also, let us call any joint state-action distribution $\mu \in \mathcal{P}(S \times X)$ an in-action environment, because the one formed by other players is what a player sees in the midst of play in any period.

4 The Nonatomic Game

Given an initial pre-action environment $\sigma_1 \in \mathcal{P}(S)$, we can define a nonatomic game $\Gamma(\sigma_1)$ which starts period 1 with σ_1 as the distribution of all players' states. For it, we can use $x_{[1\bar{t}]} \equiv (x_t \mid t = 1, ..., \bar{t}) \in (\mathcal{M}(S \times G, X))^{\bar{t}}$ to denote a policy profile. Here, each $x_t \in \mathcal{M}(S \times G, X)$ is a map from a player's state-shock pair to his prescribed action to take. Along with the given initial environment σ_1, this profile will help generate a deterministic pre-action environment trajectory $\sigma_{[1,\bar{t}+1]} \equiv (\sigma_t \mid t = 1, 2, ..., \bar{t}, \bar{t}+1) \in (\mathcal{P}(S))^{\bar{t}+1}$. This allows a player's policy to be observation-blind; that is, what portion of σ_t is observable to the player in each period t is not of any concern. The determinism of the environment evolution in $\Gamma(\sigma_1)$ can be heuristically argued along the line of LLN. However, in view of Feldman and Giles [10] and Judd [19], a rigorous proof does not seem to be in the offing. Thus, we have made this feature an axiom of the nonatomic game rather than a property derivable from its finite counterparts.

In the following, we discuss how the deterministic trajectory can be formed. Let $t = 1, ..., \bar{t}$ be given. When all players form state distribution $\sigma_t \in \mathcal{P}(S)$ at the beginning and adopt the same plan $x_t \in \mathcal{M}(S \times G, X)$, the in-action environment $\mu_t \in \mathcal{P}(S \times X)$ to be experienced by all players will take the form

$$\mu_t = (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, \tag{9}$$

where 1_S stands for the projection map from $S \times G$ to S. The meaning for (9) is that, for any measurable joint state-action set $W' \in \mathcal{B}(S \times X)$,

$$\mu_t(W') = (\sigma_t \times \bar{\gamma})((1_S, x_t)^{-1}(W')) = \int_S \int_G \mathbf{1}[(s, x_t(s, g)) \in W'] \cdot \bar{\gamma}(dg) \cdot \sigma_t(ds). \tag{10}$$

This reflects that the joint distribution for states and action-inducing shocks is the product form $\sigma_t \times \bar{\gamma}$; also, x_t provides the map from state-shock pairs to actions for this period.

For a player who starts with state s_t and has experienced action-inducing shock g_t as well as transition-inducing shock i_t, his new state will be governed by (2) to assume

$$s_{t+1} = \tilde{s}_t(s_t, x_t(s_t, g_t), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i_t). \tag{11}$$

To describe the transition of the overall pre-action environment from σ_t to σ_{t+1} under action plan x_t, we define operator $T_t(x_t)$ on $\mathcal{P}(S)$. Note that states are distributed according to σ_t, action-inducing shocks are distributed according to $\bar{\gamma}$, and transition-inducing shocks are

distributed according to $\bar{\iota}$. So following (11),

$$\sigma_{t+1} = T_t(x_t) \circ \sigma_t = (\sigma_t \times \bar{\iota} \times \bar{\gamma}) \cdot [\tilde{s}_t(1_S, x_t \cdot 1_{S\times G}, (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, 1_I)]^{-1}, \tag{12}$$

meaning that, for any measurable action set $S' \in \mathcal{B}(S)$,

$$\begin{aligned} \sigma_{t+1}(S') &= [T_t(x_t) \circ \sigma_t](S') \\ &= \textstyle\int_S \int_G \int_I \mathbf{1}\left[\tilde{s}_t(s, x_t(s,g), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i) \in S'\right] \cdot \bar{\iota}(di) \cdot \bar{\gamma}(dg) \cdot \sigma_t(ds). \end{aligned} \tag{13}$$

We can iteratively define $T_{[tt']}(x_{[tt']})$ for $t' = t-1, t, t+1, ...$ so that $T_{[t,t-1]}$ is the identity mapping on $\mathcal{P}(S)$ and for $t' = t, t+1, ...,$

$$T_{[tt']}(x_{[tt']}) = T_{t'}(x_{t'}) \circ T_{[t,t'-1]}(x_{[t,t'-1]}). \tag{14}$$

The environment trajectory alluded to earlier is therefore

$$\sigma_{[1,\bar{t}+1]} = (T_{[1,t-1]}(x_{[1,t-1]}) \circ \sigma_1 \mid t = 1, 2, ..., \bar{t}, \bar{t}+1). \tag{15}$$

In defining $\Gamma(\sigma_1)$'s equilibria, we subject a candidate policy profile to the one-time deviation of a single player, who is negligible in his influence over others. The deviation will not alter the environment trajectory corresponding to the candidate profile. Thus, we define $v_t(s_t, \sigma_t, x_{[t\bar{t}]}, y_t)$ as the total expected payoff a player can make from time t to $\bar{t}$, when he starts with state $s_t \in S$, other players form pre-action environment $\sigma_t \in \mathcal{P}(S)$, all players adopt policy $x_{[t\bar{t}]} \equiv (x_t, ..., x_{\bar{t}}) \in (\mathcal{M}(S \times G, X))^{\bar{t}-t+1}$ with the exception of the current player in period t alone, who then deviates to policy $y_t \in \mathcal{M}(S \times G, X)$ in that period. As a terminal condition, we have

$$v_{\bar{t}+1}(s_{\bar{t}+1}, \sigma_{\bar{t}+1}, y_{\bar{t}+1}) = 0. \tag{16}$$

For $t = \bar{t}, \bar{t}-1, ..., 1$, we have

$$\begin{aligned} v_t(s_t, \sigma_t, x_{[t\bar{t}]}, y_t) = \textstyle\int_G [\tilde{f}_t(s_t, y_t(s_t, g_t), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}) + \int_I v_{t+1}(\tilde{s}_t(s_t, y_t(s_t, g_t), \\ (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i_t), T_t(x_t) \circ \sigma_t, x_{[t+1,\bar{t}]}, x_{t+1}) \cdot \bar{\iota}(di_t)] \cdot \bar{\gamma}(dg_t), \end{aligned} \tag{17}$$

due to the dynamics illustrated in (11) to (13). The deviation y_t affects the current player's action in period t and his own state in period $t+1$. But as a distinctive feature of the NG setup, it has no bearing on the period-$(t+1)$ pre-action environment $T_t(x_t) \circ \sigma_t$.

Now define $u_t : \mathcal{P}(S) \times (\mathcal{M}(S \times G, X))^{\bar{t}-t+1} \times \mathcal{M}(S \times G, X) \longrightarrow \Re$ so that

$$u_t(\sigma_t, x_{[t\bar{t}]}, y_t) = \int_S v_t(s_t, \sigma_t, x_{[t\bar{t}]}, y_t) \cdot \sigma_t(ds_t). \tag{18}$$

This can be understood as one particular player's average gain from period t onward when the same conditions specified earlier prevail and his period-t state is sampled from the distribution σ_t. We deem policy $x^*_{[1\bar{t}]} \equiv (x^*_t \mid t = 1, 2, ..., \bar{t}) \in (\mathcal{M}(S \times G, X))^{\bar{t}}$ a Markov equilibrium for the game $\Gamma(\sigma_1)$ when, for every $t = 1, 2, ..., \bar{t}$ and $y_t \in \mathcal{M}(S \times G, X)$,

$$u_t(T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1, x^*_{[t\bar{t}]}, x^*_t) \geq u_t(T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1, x^*_{[t\bar{t}]}, y_t). \tag{19}$$

That is, policy $x^*_{[1\bar{t}]}$ will be regarded an equilibrium when it cannot be beaten by any plan $y_t \in \mathcal{M}(S \times G, X)$ in any period t in an average sense that is defined by the period-t environment $\sigma_t = T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1$. We caution that (19) is weaker than

$$v_t(s_t, T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1, x^*_{[t\bar{t}]}, x^*_t) \geq v_t(s_t, T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1, x^*_{[t\bar{t}]}, y_t), \tag{20}$$

for every $s_t \in S$. On the other hand, since $y_t \in \mathcal{M}(S \times G, X)$ allows for much freedom in choosing for each state $s \in S$ and shock $g \in G$ a competitive reaction $y_t(s, g)$, there is not much difference between the two criteria aside from measurability subtleties.

5 Finite-player Games

We need more notation to appropriately describe finite games. Let A be a metric space. For $a \in A$, we use ε_a to denote the singleton probability measure with $\varepsilon_a(\{a\}) = 1$. For $a = (a_1, ..., a_n) \in A^n$ where $n \in N$, the set of natural numbers, we use ε_a for $\sum_{m=1}^n \varepsilon_{a_m}/n$. The two uses are consistent. We also use $\mathcal{P}_n(A)$ for the space of probability measures of the type ε_a for $a \in A^n$, i.e., the space of empirical distributions generated from n samples.

For some $n = 2, 3, ...$ and initial state distribution $\sigma_1 \in \mathcal{P}_n(S)$, we can define an n-player game $\Gamma_n(\sigma_1)$. Note the initial pre-action environment σ_1 must be of the form $\varepsilon_{s_1} = \varepsilon_{(s_{11}, s_{12}, ..., s_{1n})}$, where each $s_{1m} \in S$ is player m's initial state. The game's payoffs and state transitions are still governed by (1) and (2), respectively. In period t, the pre-action environment is also some $\sigma_t = \varepsilon_{(s_{t1}, ..., s_{tn})} \in \mathcal{P}_n(S) \subset \mathcal{P}(S)$. Hence, the in-action environment $\mu_{t1} \in \mathcal{P}_{n-1}(S \times X) \subset \mathcal{P}(S \times X)$ experienced by any designated player 1 is the empirical distribution $\varepsilon_{s_{t,-1}, y_{t,-1}} = \varepsilon_{((s_{t2}, y_{t2}), ..., (s_{tn}, y_{tn}))}$ when each player m starts with state $s_{tm} \in S$ and takes action $y_{tm} \in X$. Let players still adopt policy $x_{[1\bar{t}]} \equiv (x_t \mid t = 1, ..., \bar{t}) \in (\mathcal{M}(S \times G, X))^{\bar{t}}$, which is but the crudest of many choices presented to the n players. We shall see later that this restriction is not going to do much harm.

Simplistic as it may seem, x will not help generate a deterministic environment trajectory. For shock vectors $g_t = (g_{t1}, ..., g_{tn}) \in G^n$ and $i_t = (i_{t1}, ..., i_{tn}) \in I^n$, we can define $T_{nt}(x_t, g_t, i_t)$

as the operator on $\mathcal{P}_n(S)$ that converts a period-t pre-action environment into a period-$(t+1)$ one. Thus following (11), $\varepsilon_{s_{t+1}} = T_{nt}(x_t, g_t, i_t) \circ \varepsilon_{s_t}$ is such that

$$s_{t+1,m} = \tilde{s}_t(s_{tm}, x_t(s_{tm}, g_{tm}), \varepsilon_{s_{t,-m},g_{t,-m}} \cdot (1_S, x_t)^{-1}, i_{tm}), \qquad \forall m = 1, 2, ..., n, \tag{21}$$

where each $\varepsilon_{s_{t,-m},g_{t,-m}}$ represents the empirical distribution built on state-shock pairs (s_{t1}, g_{t1}), ..., $(s_{t,m-1}, g_{t,m-1})$, $(s_{t,m+1}, g_{t,m+1})$, ..., (s_{tn}, g_{tn}). This reflects that player m's in-action environment is made up of the states and actions of the other $n-1$ players. Again, we define $T_{n,[tt']}$ as the identity map when $t' \leq t-1$ and when $t \leq t'$, let

$$T_{n,[tt']}(x_{[tt']}, g_{[tt']}, i_{[tt']}) = T_{nt'}(x_{t'}, g_{t'}, i_{t'}) \circ T_{n,[t,t'-1]}(x_{[t,t'-1]}, g_{[t,t'-1]}, i_{[t,t'-1]}). \tag{22}$$

The evolution of pre-action envirnoments $\sigma_t = \varepsilon_{s_t}$ is guided by the random shocks g_t and i_t, and hence is stochastic by nature.

For an n-player game, let $v_{nt}(s_{t1}, \varepsilon_{s_{t,-1}}, x_{[t\bar{t}]}, y_t)$ be the total expected payoff player 1 can make from t to $\bar{t}$, when he starts with state $s_{t1} \in S$, other players' initial environments are describable by their aggregate empirical state distribution $\varepsilon_{s_{t,-1}} = \varepsilon_{(s_{t2},...,s_{tn})}$, and all players adopt the policy $x_{[t\bar{t}]} \equiv (x_t, x_{t+1}, ..., x_{\bar{t}}) \in (\mathcal{M}(S \times G, X))^{\bar{t}-t+1}$ from period t to period $\bar{t}$ with the exception of player 1 in period t alone, who then deviates to policy $y_t \in \mathcal{M}(S \times G, X)$. As a terminal condition, we have

$$v_{n,\bar{t}+1}(s_{\bar{t}+1,1}, \varepsilon_{s_{\bar{t}+1,-1}}, y_{\bar{t}+1}) = 0. \tag{23}$$

For $t = \bar{t}, \bar{t}-1, ..., 1$, we have the recursive relationship

$$\begin{aligned} v_{nt}(s_{t1}, \varepsilon_{s_{t,-1}}, x_{[t\bar{t}]}, y_t) = \textstyle\int_{G^n} \bar{\gamma}^n(dg_t) \times \{\tilde{f}_t(s_{t1}, y_t(s_{t1}, g_{t1}), \varepsilon_{s_{t,-1},g_{t,-1}} \cdot (1_S, x_t)^{-1}) \\ + \textstyle\int_{I^n} \bar{\iota}^n(di_t) \times v_{n,t+1}(\tilde{s}_t(s_{t1}, y_t(s_{t1}, g_{t1}), \varepsilon_{s_{t,-1},g_{t,-1}} \cdot (1_S, x_t)^{-1}, i_{t1}), \\ [T_{nt}(x_t, g_t, i_t) \circ \varepsilon_{s_t}]_{-1}, x_{[t+1,\bar{t}]}, x_{t+1})\}, \end{aligned} \tag{24}$$

due to the dynamics illustrated in (11) and (21). By $[T_{nt}(x_t, g_t, i_t) \circ \varepsilon_{s_t}]_{-1}$, we mean $\varepsilon_{s_{t+1,-1}}$, where $\varepsilon_{s_{t+1}}$ is $T_{nt}(x_t, g_t, i_t) \circ \varepsilon_{s_t}$ as defined through (21). The current (24) is much more complicated than the NG counterpart (17). The evolution from period t to $t+1$ now depends on action-inducing shocks $g_t \equiv (g_{t1}, ..., g_{tn})$ and transition-inducing shocks $i_t \equiv (i_{t1}, ..., i_{tn})$. Also, the in-action environment $\varepsilon_{s_{t,-1},g_{t,-1}} \cdot (1_S, x_t)^{-1}$ experienced by player 1 excludes his own state and action, and hence is different from the environment faced by any other player. Similarly, the in-action environment $[T_{nt}(x_t, g_t, i_t) \circ \varepsilon_{s_t}]_{-1}$ to be faced by player 1 in period $t+1$ is unique to him as well. The added complexity actually helps to motivate our exploitation of the easier-to-handle NG case.

Let $\sigma_{[1\bar{t}]} \equiv (\sigma_t \mid t = 1, ..., \bar{t}) \in (\mathcal{P}(S))^{\bar{t}}$ be a sequence of environments. For $\epsilon \geq 0$, we deem $x^*_{[1\bar{t}]} \equiv (x^*_t \mid t = 1, ..., \bar{t}) \in (\mathcal{M}(S \times G, X))^{\bar{t}}$ an ϵ-Markov equilibrium for the game family $(\Gamma_n(\varepsilon_{s_1}) \mid s_1 \in S^n)$ in the sense of $\sigma_{[1\bar{t}]}$ when, for every $t = 1, ..., \bar{t}$ and $y_t \in \mathcal{M}(S \times G, X)$,

$$\int_{S^n} v_{nt}(s_{t1}, \varepsilon_{s_{t,-1}}, x^*_{[t\bar{t}]}, x^*_t) \cdot {\sigma_t}^n(ds_t) \geq \int_{S^n} v_{nt}(s_{t1}, \varepsilon_{s_{t,-1}}, x^*_{[t\bar{t}]}, y_t) \cdot {\sigma_t}^n(ds_t) - \epsilon. \tag{25}$$

That is, action plan $x^*_{[1\bar{t}]}$ will be an ϵ-Markov equilibrium in the sense of $\sigma_{[1\bar{t}]}$ when under its guidance, the average payoff from any period t on will not be improved by more than ϵ through any deviation, where the "average" is taken with respect to state distribution σ_t.

6 Main Convergence Results

We can achieve convergences of environments and then of equilibria. The former is more fundamental and challenging, and the latter is built on it.

6.1 Convergence of Environments

Even without touching upon payoffs or equilibria, we can establish a link between finite games and their NG counterpart. It reflects that stochastic environment pathways experienced by large finite games converge to the NG's deterministic environment trajectory.

Let A, B, and C be metric spaces and $\pi_B \in \mathcal{P}(B)$ be a distribution. We use $\mathcal{K}(A, B, \pi_B, C) \subseteq \mathcal{M}(A \times B, C)$ to represent the space of all measurable functions from $A \times B$ to C that are uniformly continuous in a probabilistic sense. The criterion for $y \in \mathcal{K}(A, B, \pi_B, C)$ is as follows. For any $\epsilon > 0$, there exists $\delta > 0$, so that for any $a, a' \in A$ satisfying $d_A(a, a') < \delta$,

$$\pi_B(\{b \in B | d_C(y(a, b), y(a', b)) < \epsilon\}) > 1 - \epsilon. \tag{26}$$

When B is a singleton and hence π_B is degenerate, $y \in \mathcal{K}(A, B, \pi_B, C)$ merely means that y is a uniformly continuous function from A to C, a situation we denote by $y \in \mathcal{K}(A, C)$. For regular B and π_B, the meaning is somehow that continuity will happen in most cases. Probabilistic continuity in the sense of $\mathcal{K}(A, B, \pi_B, C)$ has another much anticipated link with ordinary continuity.

Proposition 1 *(I) Let $\pi_B \in \mathcal{P}(B)$ and $y \in \mathcal{M}(A \times B, C)$. Suppose $\theta : A \longrightarrow \mathcal{P}(C)$ is crafted from*

$$\theta(a) = \pi_B \cdot (y(a, \cdot))^{-1}.$$

Then, we will have $\theta \in \mathcal{K}(A, \mathcal{P}(C))$ when $y \in \mathcal{K}(A, B, \pi_B, C)$.

(II) On the other hand, let $B = [0, 1]$, $\pi_B \in \mathcal{P}(B)$ be uniform, C be a closed subset of the real line, and θ be a mapping from A to $\mathcal{P}(C)$. Suppose $y : A \times B \longrightarrow C$ is crafted from

$$y(a, b) = \inf\{c \in C | \theta(a|(-\infty, c] \cap C) \geq b\}.$$

Then, we still have $\theta(a) = \pi_B \cdot (y(a, \cdot))^{-1}$. Suppose in addition that $y(a, \cdot)$ is continuous at an a-independent rate. Then, when $\theta(\cdot)$ is continuous at a given $a \in A$, it will follow that $y(a_n, \cdot)$ converges to $y(a, \cdot)$ in probability for any sequence a_n converging to a.

We may see from Proposition 1 that, as a requirement, $y \in \mathcal{K}(A, B, \pi_B, C)$ is stronger than $\theta(a) = \pi_B \cdot (y(a, \cdot))^{-1}$ being uniformly continuous in a, but probably not by much. Technical developments of this subsection are left to Appendix A. We now make two assumptions on the transition function $\tilde{s}_t$:

(S1) For every $\mu \in \mathcal{P}(S \times X)$, the function $\tilde{s}_t(\cdot, \cdot, \mu, \cdot)$ is a member of $\mathcal{K}(S \times X, I, \bar{\iota}, S)$. That is, for any $\mu \in \mathcal{P}(S \times X)$ and $\epsilon > 0$, there exist $\delta_S > 0$ and $\delta_X > 0$, so that for any $s, s' \in S$ and $x, x' \in X$ satisfying $d_S(s, s') < \delta_S$ and $d_X(x, x') < \delta_X$,

$$\bar{\iota}(\{i \in I \mid d_S(\tilde{s}_t(s, x, \mu, i), \tilde{s}_t(s', x', \mu, i)) < \epsilon\}) > 1 - \epsilon.$$

(S2) Not only is it true that $\tilde{s}_t(s, x, \cdot, \cdot) \in \mathcal{K}(\mathcal{P}(S \times X), I, \bar{\iota}, S)$ at every $(s, x) \in S \times X$, but the continuity is also achieved at a rate independent of the (s, x) present. That is, for any $\mu \in \mathcal{P}(S \times X)$ and $\epsilon > 0$, there exists $\delta > 0$, so that for any $\mu' \in \mathcal{P}(S \times X)$ satisfying $\rho_{S \times X}(\mu, \mu') < \delta$, as well as any $s \in S$ and $x \in X$,

$$\bar{\iota}(\{i \in I \mid d_S(\tilde{s}_t(s, x, \mu, i), \tilde{s}_t(s, x, \mu', i)) < \epsilon\}) > 1 - \epsilon.$$

According to the comment right after Proposition 1, (S1) and (S2) together are roughly about the continuity in (s, x) and in μ of the state transition kernel $\tilde{g}_t$ that is defined in (3). From the detailed study of the dynamic pricing game in Section 8, we shall see that these requirements are easy to meet.

For separable metric space A, we use $(A^n, \mathcal{B}^n(A))$ to denote the product measurable space that houses n-long sample sequences. Given $\pi \in \mathcal{P}(A)$, we use π^n to denote the product measure on $(A^n, \mathcal{B}^n(A))$. We have a technical result showing that a one-step evolution in a big game is not that much different from that in a nonatomic game.

Proposition 2 *Given separable metric space A, distribution $\pi \in \mathcal{P}(A)$, and pre-action environment $\sigma \in \mathcal{P}(S)$, suppose $s_n = (s_n(a) \mid a \in A^n)$ for each $n \in N$ is a member of $\mathcal{M}(A^n, S^n)$, and $\varepsilon_{s_n(a)}$ converges to σ in probability, to the effect that*

$$\pi^n(\{a \in A^n \mid \rho_S(\varepsilon_{s_n(a)}, \sigma) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough. Then, any $T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}$ will converge to $T_t(x) \circ \sigma$ in probability for any probabilistically continuous x. That is, for any $x \in \mathcal{K}(S, G, \bar{\gamma}, X)$,

$$(\pi \times \bar{\gamma} \times \bar{\iota})^n(\{(a, g, i) \in (A \times G \times I)^n \mid \rho_S(T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}, T_t(x) \circ \sigma) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough.

Recall that ρ_S is the Prohorove metric for measuring the distance between two state distributions. Also, the operator $T_t(x)$ delineating the period-t transition of an NG's pre-action environment is defined at (13), and its finite-game counterpart $T_{nt}(x, g, i)$ is defined at (21). The proof of Proposition 2 calls upon Lemma 3 in Appendix A. This is why the spaces S, G, and I are required to be complete. Now imagine that $(A, \mathcal{B}(A), \pi)$ provides exogenous shocks that drive games' evolutions up to period t: $A = S \times G^{t-1} \times I^{t-1}$ and $\pi = \sigma_1 \times \bar{\gamma}^{t-1} \times \bar{\iota}^{t-1}$. Proposition 2 states that, when starting period t with initial state vectors $s_n(a)$ in n-player games that in aggregation increasingly resemble the given starting distribution σ for the NG, one will still get state vectors in large games that in aggregation resemble the NG's state distribution after the period-t transition. When exploiting this proposition iteratively, we can arrive at our first main result on the convergence of environments.

Theorem 1 *Let a policy profile $x_{[t\bar{t}]} \in (\mathcal{M}(S \times G, X))^{\bar{t}-t+1}$ for periods $t, t+1, ..., \bar{t}$ be such that each $x_{t'}$ is a member of $\mathcal{K}(S, G, \bar{\gamma}, X)$. Then, when we sample $s_t = (s_{t1}, ..., s_{tn})$ from a given pre-action environment $\sigma_t \in \mathcal{P}(S)$, the sequence $(T_{n,[t,t'-1]}(x_{[t,t'-1]}, g_{[t,t'-1]}, i_{[t,t'-1]}) \circ \varepsilon_{s_t}) \mid t' = t, t+1, ..., \bar{t}, \bar{t}+1)$ will converge to $(T_{[t,t'-1]}(x_{[t,t'-1]}) \circ \sigma_t \mid t' = t, t+1, ..., \bar{t}, \bar{t}+1)$ in probability. That is, for any $\epsilon > 0$ and any n large enough,*

$$(\sigma_t \times \bar{\gamma}^{\bar{t}-t+1} \times \bar{\iota}^{\bar{t}-t+1})^n(\tilde{A}_n(\epsilon)) > 1 - \epsilon,$$

where $\tilde{A}_n(\epsilon) \in \mathcal{B}^n(S \times G^{\bar{t}-t+1} \times I^{\bar{t}-t+1})$ is such that, for any $(s_t, g_{[t,\bar{t}]}, i_{[t,\bar{t}]}) \in \tilde{A}_n(\epsilon)$,

$$\rho_S(T_{n,[t,t'-1]}(x_{[t,t'-1]}, g_{[t,t'-1]}, i_{[t,t'-1]}) \circ \varepsilon_{s_t}, T_{[t,t'-1]}(x_{[t,t'-1]}) \circ \sigma_t) < \epsilon, \quad \forall t' = t, t+1, ..., \bar{t}, \bar{t}+1.$$

The multi-period transition operator $T_{[t,t'-1]}(x_{[t,t'-1]})$ for the NG is defined at (14), and its finite-game counterpart $T_{n,[t,t'-1]}(x_{[t,t'-1]}, g_{[t,t'-1]}, i_{[t,t'-1]})$ is defined at (22). Suppose an NG starts period t with pre-action environment σ_t and a slew of finite games start the period with pre-action environments that are sampled from σ_t. Let the evolution of both types of games be guided by players acting according to the same probabilistically continuous policy profile $x_{[t\bar{t}]}$. Then, as the numbers of players n involved in finite games grow to infinity, Theorem 1 predicts for ever less chances for the finite games' period-t' environments $T_{n,[t,t'-1]}(x_{[t,t'-1]}, g_{[t,t'-1]}, i_{[t,t'-1]}) \circ \varepsilon_{s_t}$ to be even slightly away from the NG's deterministic period-t' environment $T_{[t,t'-1]}(x_{[t,t'-1]}) \circ \sigma_t$. In view of the comment right after Proposition 1, the requirement that each $x_{t'} \in \mathcal{K}(S, G, \bar{\gamma}, X)$ translates roughly into the continuity in s of the corresponding random action plan $\chi_{t'}$ that is defined in (6).

6.2 Convergence of Equilibria

We now set out to establish this section's main result, that an equilibrium from the NG will serve as an ever more accurate approximate equilibrium for ever larger finite games. First, we need to assume that the single-period payoff functions $\tilde{f}_t$ are continuous:

(F1) Each $\tilde{f}_t(s, x, \mu)$ is continuous in (s, x).

(F2) Each $\tilde{f}_t(s, x, \mu)$ is continuous in μ at an (s, x)-independent rate.

There are a couple of intermediate results, whose proofs are provided in Appendix B.

Proposition 3 *$v_t(s_t, \sigma_t, x_{[t\bar{t}]}, x_t)$ is continuous in s_t under probabilistically continuous $x_{t'}$'s.*

Proposition 4 *Let $\sigma_t \in \mathcal{P}(S)$ and $x_{[t\bar{t}]} \in (\mathcal{K}(S, G, \bar{\gamma}, X))^{\bar{t}-t+1}$ be given. Suppose sequence $s_{t,-1} = (s_{t2}, s_{t3}, ...)$ is sampled from σ_t, then $v_{nt}(s_{t1}, \varepsilon_{s^n_{t,-1}}, x_{[t\bar{t}]}, x_t)$ will converge to $v_t(s_{t1}, \sigma_t, x_{[t\bar{t}]}, x_t)$ in probability at an s_{t1}-independent rate, where $s^n_{t,-1}$ stands for the cutoff $(s_{t2}, s_{t3}, ..., s_{tn})$.*

Now here comes our main transient result.

Theorem 2 *For some state distribution $\sigma_1 \in \mathcal{P}(S)$, suppose $x^*_{[1\bar{t}]} \equiv (x^*_t \mid t = 1, 2, ..., \bar{t}) \in (\mathcal{K}(S, G, \bar{\gamma}, X))^{\bar{t}}$ is a probabilistically continuous Markov equilibrium of the nonatomic game $\Gamma(\sigma_1)$. Then, for any $\epsilon > 0$ and large enough $n \in N$, this $x^*_{[1\bar{t}]}$ is also an ϵ-Markov equilibrium for the game family $(\Gamma_n(\varepsilon_{s_1}) \mid s_1 \in S^n)$ in the sense of $(T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1 \mid t = 1, ..., \bar{t})$, meaning that for any $t = 1, ..., \bar{t}$ and $y_t \in \mathcal{M}(S \times G, X)$, (25) is true with*

$$\sigma_t = T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1.$$

Furthermore, the same is true in the sense of $(T_{n,[1,t-1]}(x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]})\circ\sigma_1 \mid t = 1, ..., \bar{t})$ *in probability. The latter means that, for any* $\epsilon > 0$ *and large enough* $n \in N$*, for any* $t = 1, ..., \bar{t}$ *and* $y_t \in \mathcal{M}(S \times G, X)$*,*

$$\begin{aligned}
&\textstyle\int_{S^n} \sigma_1^n(ds_1) \times \int_{G^{n\cdot(t-1)}} \bar{\gamma}^{n\cdot(t-1)}(dg_{1,t-1]}) \times \int_{I^{n\cdot(t-1)}} \bar{\iota}^{n\cdot(t-1)}(di_{[1,t-1]})\times \\
&\qquad \times v_{nt}(s_{t,1}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}), \varepsilon_{s_{t,-1}}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}), x^*_{[t\bar{t}]}, x^*_t) \\
&\geq \textstyle\int_{S^n} \sigma_1^n(ds_1) \times \int_{G^{n\cdot(t-1)}} \bar{\gamma}^{n\cdot(t-1)}(dg_{1,t-1]}) \times \int_{I^{n\cdot(t-1)}} \bar{\iota}^{n\cdot(t-1)}(di_{[1,t-1]})\times \\
&\qquad \times v_{nt}(s_{t,1}(s_1, x^*_{1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}), \varepsilon_{s_{t,-1}}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}), x^*_{[t\bar{t}]}, y_t) - \epsilon,
\end{aligned}$$

where both $s_{t,1}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]})$ *and* $\varepsilon_{s_{t,-1}}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]})$ *should be understood as being extracted from*

$$\varepsilon_{s_t}(s_1, x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}) = T_{n,[1,t-1]}(x^*_{[1,t-1]}, g_{[1,t-1]}, i_{[1,t-1]}) \circ \varepsilon_{s_1}.$$

Theorem 2 says that, when there are enough of them, players in a finite game can agree on an NG equilibrium and expect to lose little on average; also, the distribution based on which "average" is taken is itself an accurate assessment of what players' states would be had they followed the NG equilibrium all along. The two versions used here have their counterparts in the first two interpretations of Yang's [38] Theorem 2, concerning discrete state and action spaces. In the NG limit, the evolution of pre-action environments is deterministic. An equilibrium here, which is necessarily observation-blind, serves as a good asymptotic equilibrium for finite games when there are enough players; and, this asymptotic result is independent of the observatory power of players in the finite games. The root cause of this is, as stated in Theorem 1, that the aggregate environment in large games evolves in a nearly deterministic fashion and hence the value of observation has vanished.

7 A Stationary Situation

Now we study an infinite-horizon model with stationary features. To this end, we suppose there is a state transition function $\tilde{s}$, so that

$$\tilde{s}_t(s, x, \mu, i) = \tilde{s}(s, x, \mu, i), \qquad \forall t = 1, 2, ... \tag{27}$$

In addition, we suppose there is a payoff function $\tilde{f}$, so that

$$\tilde{f}_t(s, x, \mu) = \bar{\phi}^{\,t-1} \cdot \tilde{f}(s, x, \mu), \qquad \forall t = 1, 2, ..., \tag{28}$$

where $\bar{\phi} \in [0,1)$ is a discount factor. Also, we use $\bar{f}$ for the bound $\bar{f}_1$ that appears in (1).

For the nonatomic game Γ with the above stationary features, we use $x \equiv (x(s,g) \mid s \in S, g \in G) \in \mathcal{M}(S \times G, X)$ to represent a stationary policy profile. It is a map from the current period's state and action-inducing shock to the player's action. Given an $x \in \mathcal{M}(S \times G, X)$, we denote by $T(x)$ the operator on $\mathcal{P}(S)$ that converts one state distribution σ to another σ' so that, for every $S' \in \mathcal{B}(S)$,

$$\begin{aligned} \sigma'(S') &= [T(x) \circ \sigma](S') \\ &= \textstyle\int_S \int_G \int_I \mathbf{1}(\tilde{s}(s, x(s,g), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, i) \in S') \cdot \bar{\iota}(di) \cdot \bar{\gamma}(dg) \cdot \sigma(ds). \end{aligned} \tag{29}$$

An environment $\sigma \in \mathcal{P}(S)$ is said to be associated with x when

$$\sigma = T(x) \circ \sigma. \tag{30}$$

That is, we consider $\sigma \in \mathcal{P}(S)$ to be associated with $x \in \mathcal{M}(S \times G, X)$ when the former is invariant under the state transition facilitated by the $T(x)$ operator.

Suppose stationary policy $x \in \mathcal{M}(S \times G, X)$ is associated with pre-action environment $\sigma \in \mathcal{P}(S)$. For $t = 0, 1, ...$, we define $v_t(s, \sigma, x, y)$ as the total expected payoff a player can make from period 1 to t, when he starts period 1 with state $s \in S$ and outside environment $\sigma \in \mathcal{P}(S)$, while all players keep on using policy x from period 1 to t with the exception of the current player in the very beginning, who then deviates to $y \in \mathcal{M}(S \times G, X)$. As a terminal condition, we have

$$v_0(s, \sigma, x, y) = 0. \tag{31}$$

Due to the stationarity of the setting, we have, for $t = 1, 2, ...$,

$$\begin{aligned} v_t(s, \sigma, x, y) = \textstyle\int_G [\tilde{f}(s, y(s,g), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}) \\ + \bar{\phi} \cdot \textstyle\int_I v_{t-1}(\tilde{s}(s, y(s,g), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, i), \sigma, x, x) \cdot \bar{\iota}(di)] \cdot \bar{\gamma}(dg). \end{aligned} \tag{32}$$

Using (31) and (32), we can inductively show that

$$\mid v_{t+1}(s, \sigma, x, y) - v_t(s, \sigma, x, y) \mid \leq \bar{\phi}^t \cdot \bar{f}. \tag{33}$$

The sequence $\{v_t(s, \sigma, x, y) \mid t = 0, 1, ...\}$ is thus Cauchy with a limit point $v_\infty(s, \sigma, x, y)$. This $v_\infty(s, \sigma, x, y)$ can be understood as the infinite-horizon total discounted expected payoff a player can obtain by starting with state s and environment σ, while all players adhere to the action plan x except for the current player in the beginning, who deviates to y then.

We deem $x^* \in \mathcal{M}(S \times G, X)$ a Markov equilibrium for the nonatomic game Γ when for some $\sigma^* \in \mathcal{P}(S)$ associated with x^* in the fashion of (30) and every $y \in \mathcal{M}(S \times G, X)$,

$$\int_S v_\infty(s, \sigma^*, x^*, x^*) \cdot \sigma^*(ds) \geq \int_S v_\infty(s, \sigma^*, x^*, y) \cdot \sigma^*(ds). \tag{34}$$

Therefore, a policy will be considered an equilibrium when it induces an invariant environment profile under which the policy forms a best response in the long run.

Now we move on to the n-player game Γ_n with the same stationary features provided by $\bar{s}$, $\bar{f}$, and $\bar{\phi}$. Given policy profile $x = (x(s,g) \mid s \in S, g \in G) \in \mathcal{M}(S \times G, X)$, action-inducing shock vector $g = (g_1, ..., g_n) \in G^n$, and transition-inducing shock vector $i = (i_1, ..., i_n) \in I^n$, we define $T_n(x, g, i)$ as the operator on $\mathcal{P}_n(S)$ that converts a period's pre-action environment into that of a next period. Thus, $\varepsilon_{s'} = T_n(x, g, i) \circ \varepsilon_s$ is such that

$$s'_m = \tilde{s}(s_m, x(s_m, g_m), \varepsilon_{s_{-m}, g_{-m}} \cdot (1_S, x)^{-1}, i_m), \qquad \forall m = 1, 2, ..., n. \tag{35}$$

Let $v_{nt}(s_1, \varepsilon_{s_{-1}}, x, y)$ be the total expected payoff player 1 can make from period 1 to t, when the player's starting state is $s_1 \in S$, other players' initial environments is describable by their aggregate empirical state distribution $\varepsilon_{s_{-1}} = \varepsilon_{(s_2, ..., s_n)}$, and all players adopt the policy $x \in \mathcal{M}(S \times G, X)$ with the exception that player 1 adopts policy $y \in \mathcal{M}(S \times G, X)$ in the very beginning. As a terminal condition, we have

$$v_{n0}(s_1, \varepsilon_{s_{-1}}, x, y) = 0. \tag{36}$$

For $t = \bar{t}, \bar{t} - 1, ..., 1$, we have the recursive relationship

$$\begin{array}{l} v_{nt}(s_1, \varepsilon_{s_{-1}}, x, y) = \int_{G^n} \bar{\gamma}^n(dg) \times \{\tilde{f}(s_1, y(s_1, g_1), \varepsilon_{s_{-1}, g_{-1}} \cdot (1_S, x)^{-1}) + \bar{\phi} \times \\ \times \int_{I^n} \bar{\iota}^n(di) \times v_{n,t-1}(\tilde{s}(s_1, y(s_1, g_1), \varepsilon_{s_{-1}, g_{-1}} \cdot (1_S, x)^{-1}, i_1), [T_n(x, g, i) \circ 1_S]_{-1}, x, x)\}, \end{array} \tag{37}$$

where $[T_n(x, g, i) \circ 1_S]_{-1}$ stands for $\varepsilon_{s'_{-1}}$, such that $\varepsilon_{s'} = T_n(x, g, i) \circ 1_S$. Using (36) and (37), we can inductively show that

$$\mid v_{n,t+1}(s_1, \varepsilon_{s_{-1}}, x, y) - v_{nt}(s_1, \varepsilon_{s_{-1}}, x, y) \mid \leq \bar{\phi}^t \cdot \bar{f}. \tag{38}$$

Thus, the sequence $\{v_{nt}(s_1, \varepsilon_{s_{-1}}, x, y) \mid t = 0, 1, ...\}$ is Cauchy with limit $v_{n\infty}(s_1, \varepsilon_{s_{-1}}, x, y)$.

We make the following assumptions, which are t-independent versions of (S1) to (F2):

(S1-s) For every $\mu \in \mathcal{P}(S \times X)$, the function $\tilde{s}(\cdot, \cdot, \mu, \cdot)$ is a member of $\mathcal{K}(S \times X, I, \bar{\iota}, S)$.

(S2-s) Not only is it true that $\tilde{s}(s, x, \cdot, \cdot) \in \mathcal{K}(\mathcal{P}(S \times X), I, \bar{\iota}, S)$ at every $(s, x) \in S \times X$, but the continuity is also achieved at a rate independent of the (s, x) present.

(F1-s) The function $\tilde{f}(s,x,\mu)$ is continuous in (s,x).

(F2-s) The function $\tilde{f}(s,x,\mu)$ is continuous in μ at an (s,x)-independent rate.

Here comes our main result for the stationary case.

Theorem 3 *Suppose $x^* \in \mathcal{K}(S,G,\bar{\gamma},X)$ is a probabilistically continuous Markov equilibrium for the nonatomic game Γ. Let $\sigma^* \in \mathcal{P}(S)$ be associated with x^* in the fashion of (30). Then, for any $\epsilon > 0$ and large enough $n \in N$, for any $y \in \mathcal{M}(S \times G, X)$,*

$$\int_{S^n} v_{n\infty}(s_1, \varepsilon_{s_{-1}}, x^*, x^*) \cdot (\sigma^*)^n(ds) \geq \int_{S^n} v_{n\infty}(s_1, \varepsilon_{s_{-1}}, x^*, y) \cdot (\sigma^*)^n(ds) - \epsilon.$$

The proof of Theorem 3 is provided in Appendix C. It states that players in a large finite game will not regret much by keeping on adopting a stationary equilibrium for its correspondent nonatomic game. The regret is measured in an average sense, where the underlying invariant state distribution for measuring "average" is part of the NG equilibrium. So players can fare well by responding to their individual states in the same fashion all the way into the infinite future.

8 A Dynamic Pricing Game

We study a dynamic pricing game participated by heterogeneous firms. Since the random demand arrival process is influenced by prices charged by all firms and leftover items are stored for future sales, the finite-player version of this problem is virtually intractable. The usefulness of the transient result Theorem 2 is thus at full display. To the stationary case also involving production, the stationary result Theorem 3 can further be applied. Moreover, depending on which portion of the outside environment, whether it be merely other firms' prices or both their prices and inventory levels, are observable, there can be different versions of the finite game. The NG approximation renders these differences irrelevant.

8.1 Model Primitives

Here, players are heterogeneous firms engaged in price competition whilst items unsold in a period are put aside for future sales. For strictly positive integers $\bar{s}_T$ and $\bar{s}_I$, let $S_T \equiv \{1, ..., \bar{s}_T\}$ be the set of firm types and $S_I \equiv \{0\} \cup \{1, ..., \bar{s}_I\}$ be the set of potential inventory levels. The product $S \equiv \bar{S}_T \times \bar{S}_I$ containing $\bar{s} \equiv \bar{s}_T \times (\bar{s}_I + 1)$ elements constitutes the

current problem's state space. We let action space be the set $X \equiv \{1, ..., \bar{x}\} \cup \{\infty\}$ of price choices, whereas the price associated with each choice x is denoted by $\bar{p}_x$. Suppose that

$$0 < \bar{p}_1 < \cdots < \bar{p}_{\bar{x}} < \bar{p}_\infty < +\infty. \tag{39}$$

Firms with no stock left and only those firms will charge the demand-stopping price $\bar{p}_\infty$.

For any natural number n, let $\Delta(n)$ be the simplex in the Euclidean space $\Re^n$:

$$\Delta(n) = \{(y_1, ..., y_n) \in [0,1]^n | y_1 + \cdots + y_n \leq 1\}, \tag{40}$$

whose boundary

$$\partial\Delta(n) = \{(y_1, ..., y_n) \in [0,1]^n | y_1 + \cdots + y_n = 1\}. \tag{41}$$

There are homeomorphisms between $\mathcal{P}(\{1, ..., n\})$ equipped with the Prohorov metric and $\Delta(n-1)$ or $\partial\Delta(n)$ equipped with any metric that is equivalent to the Euclidean one; hence, we can equate $\mathcal{P}(S) \equiv \mathcal{P}(S_T \times S_I)$ with $\partial\Delta(\bar{s}) \equiv \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$. A pre-action environment, namely, a joint type-inventory distribution $\sigma \in \mathcal{P}(S_T \times S_I)$ can be represented by a vector $(\sigma_{s_T s_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$. Within the vector, each $\sigma_{s_T s_I}$ represents the fraction of firms that are of type s_T and with s_I items left.

Besides locking the price choice ∞ with inventory level 0, we suppose that a firm's type will not affect demands going to other firms once it has run out of stock. Thus, the space of in-action environments need only be the $(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x} + 1)$-member set $(S_T \times \{1, ..., \bar{s}_I\} \times \{1, ..., \bar{x}\}) \cup \{(0, \infty)\}$ rather than the $\bar{s}_T \cdot (\bar{s}_I + 1) \cdot (\bar{x} + 1)$-member set $S_T \times S_I \times X$. In the former set, the element $(0, \infty)$ stands for the situation where the firm is of an arbitrary type, with inventory level 0 and hence price choice ∞. We can equate $\mathcal{P}((S_T \times \{1, ..., \bar{s}_I\} \times \{1, ..., \bar{x}\}) \cup \{(0, \infty)\})$ with $\Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x})$. Every joint type-inventory-price distribution $\mu \in \mathcal{P}((S_T \times \{1, ..., \bar{s}_I\} \times \{1, ..., \bar{x}\}) \cup \{(0, \infty)\})$ can be understood as a vector $(\mu_{s_T s_I x})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I, x=1,...,\bar{x}} \in \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x})$, where each $\mu_{s_T s_I x}$ stands for the fraction of firms with type s_T, inventory level s_I, and price choice x. The number $\mu_{0\infty} \equiv 1 - \sum_{s_T=1}^{\bar{s}_T} \sum_{s_I=1}^{\bar{s}_I} \sum_{x=1}^{\bar{x}} \mu_{s_T s_I x}$ will be the fraction of firms without any stock.

We further suppose that demand coming to a firm, though potentially affected by other firms' types and price choices, is unperturbed by other firms' inventory levels. It is thus necessary to deal with the space $(S_T \times \{1, ..., \bar{x}\}) \cup \{\infty\}$. The element ∞ stands for the situation where the firm is of an arbitrary type and with price choice ∞. We can equate $\mathcal{P}((S_T \times \{1, ..., \bar{x}\}) \cup \{\infty\})$ with $\Delta(\bar{s}_T \cdot \bar{x})$. A joint type-price distribution $\xi \in \mathcal{P}((S_T \times \{1, ..., \bar{x}\}) \cup \{\infty\})$ can be represented by a vector $(\xi_{s_T x})_{s_T=1,...,\bar{s}_T, x=1,...,\bar{x}} \in \Delta(\bar{s}_T \cdot \bar{x})$, such that

each $\xi_{s_T x}$ represents the fraction of firms that are of type s_T and charging price $\bar{p}_x$. Now, the value $\xi_\infty \equiv 1 - \sum_{s_T=1}^{\bar{s}_T} \sum_{x=1}^{\bar{x}} \xi_{s_T x}$ will represent the fraction of firms without any stock. The assumptions placed so far on demand arrivals are to make them closer to reality. They do not make the analysis easier. A less restrictive model where demand depends on elements of $\mathcal{P}(S \times X) \equiv \mathcal{P}(S_T \times (\{0\} \cup \{1, ..., \bar{s}_I\}) \times (\{1, ..., \bar{x}\} \cup \{\infty\}))$ can be similarly handled.

For transition-inducing shocks, we let $I = [0, 1]$ and the distribution $\bar{\iota}$ be uniform. For each $s_T = 1, ..., \bar{s}_T$ and $x = 1, ..., \bar{x}$, let $\tilde{\lambda}_{s_T x}(\cdot)$ be a function from $\Delta(\bar{x})$ to $[0, 1]$. In a period, suppose a particular type-s_T firm charges price $x = 1, ..., \bar{x}$, faces an outside type-price distribution $\xi \in \Delta(\bar{s}_T \cdot \bar{x})$, and receives a shock $i \in I$. Then, it will experience

$$\text{a unit demand when } i \leq \tilde{\lambda}_{s_T x}(\xi) \text{ and no demand otherwise.} \tag{42}$$

Therefore, demand is effectively distributed as a Bernoulli random variable with parameter $\tilde{\lambda}_{s_T x}(\xi)$ when the firm is of type s_T and charges price $\bar{p}_x$, while the outside environment induces $\xi \in \Delta(\bar{s}_T \cdot \bar{x})$. When periods are short, this setup is versatile enough to accommodate a variety of situations involving random price-demand relationships and independent demand arrivals. We make the following continuity assumption:

(P1) Every $\tilde{\lambda}_{s_T x}(\cdot)$ is continuous from $\Delta(\bar{s}_T \cdot \bar{x})$ to $[0, 1]$.

This is a mild condition. For instance, for positive constants $\bar{\alpha}_{11}, ..., \bar{\alpha}_{\bar{s}_T \bar{x}}$, $\bar{\beta}_{11}, ..., \bar{\beta}_{\bar{s}_T \bar{x}}, \bar{\beta}_\infty$, and $\bar{\gamma}_{11}, ..., \bar{\gamma}_{\bar{s}_T \bar{x}}, \bar{\gamma}_\infty$ that satisfy, for every $s_T = 1, ..., \bar{s}_T$ and $x = 1, ..., \bar{x}$,

$$\bar{\alpha}_{s_T x} \leq \frac{(\min_{s'_T=1}^{\bar{s}_T} \min_{x'=1}^{\bar{x}} \bar{\beta}_{s'_T x'} \cdot \exp(-\bar{\gamma}_{s'_T x'} \cdot \bar{p}_{x'})) \wedge (\bar{\beta}_\infty \cdot \exp(-\bar{\gamma}_\infty \cdot \bar{p}_\infty))}{\bar{\beta}_{s_T x} \cdot \exp(-\bar{\gamma}_{s_T x} \cdot \bar{p}_x)}, \tag{43}$$

it allows each $\tilde{\lambda}_{s_T x}(\xi)$ to be equal to

$$\frac{\bar{\alpha}_{s_T x} \cdot \bar{\beta}_{s_T x} \cdot \exp(-\bar{\gamma}_{s_T x} \cdot \bar{p}_x)}{\sum_{s'_T=1}^{\bar{s}_T} \sum_{x'=1}^{\bar{x}} \xi_{s'_T x'} \cdot \bar{\beta}_{s'_T x'} \cdot e^{-\bar{\gamma}_{s'_T x'} \cdot \bar{p}_{x'}} + (1 - \sum_{s'_T=1}^{\bar{s}_T} \sum_{x'=1}^{\bar{x}} \xi_{s'_T x'}) \cdot \bar{\beta}_\infty \cdot e^{-\bar{\gamma}_\infty \cdot \bar{p}_\infty}}. \tag{44}$$

But the above is consistent with a multinomial logit price-demand model. Let us also suppose that unit holding cost per period for a type-s_T firm is $\bar{h}_{s_T}$.

8.2 The Transient Game

Given a joint type-inventory-price distribution $\mu \equiv (\mu_{s_T s_I x})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I, x=1,...,\bar{x}} \in \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x})$, its type-price marginal $\mu|_{S_T \times X} \equiv (\mu|_{S_T \times X, s_T x})_{s_T=1,...,\bar{s}_T, x=1,...,\bar{x}} \in \Delta(\bar{s}_T \cdot \bar{x})$ satisfies

$$\mu|_{S_T \times X, s_T x} = \sum_{s_I=1}^{\bar{s}_I} \mu_{s_T s_I x}, \qquad \forall s_T = 1, ..., \bar{s}_T, x = 1, ..., \bar{x}. \tag{45}$$

For the single-period payoff $\tilde{f}_t$ defined in (1), we have

$$\tilde{f}_t(s_T, s_I, x, \mu) = \mathbf{1}(s_I \geq 1) \cdot \bar{p}_x \cdot \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X}) - \bar{h}_{s_T} \cdot s_I. \tag{46}$$

The above (46) reflects that a type-s_T firm can earn $\bar{p}_x$ with a $\tilde{\lambda}_{s_T x}(\mu|_{S_T \times X})$ chance when it charges the price $\bar{p}_x$ at a nonzero inventory level and when the outside environment is μ. The term $\bar{h}_{s_T} \cdot s_I$ reflects the holding cost, which we have assessed on the pre-selling inventory level. For the state-transition map $\tilde{s}_t \equiv (\tilde{s}_{T,t}, \tilde{s}_{I,t})$ as defined through (2),

$$\begin{cases} \tilde{s}_{T,t}(s_T, s_I, x, \mu, i) = s_T, \\ \tilde{s}_{I,t}(s_T, s_I, x, \mu, i) = \mathbf{1}(s_I \geq 1) \cdot [s_I - \mathbf{1}(i \leq \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X}))]. \end{cases} \tag{47}$$

The T-portion of (47) confirms that a firm's type is its innate characteristics unmovable by either time or any other factor. The I-portion reflects that, if a type-s_T firm starts a period with inventory level $s_I \geq 1$ and outside environment μ, charges price $\bar{p}_x$, and experiences shock i in the period, then it will secure a unit demand as long as $i \leq \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X})$; also, the firm's inventory will remain empty once it reaches 0.

For action-inducing shocks, we let $G = [0, 1]$ and $\bar{\gamma}$ be the uniform distribution. This is as though every firm is equipped with a random number generator and use it to assist in decision making. We focus only on action plans $x_t \in \mathcal{M}(S \times G, X)$ that correspond to members of the space $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$. Given any $\chi_t \equiv (\chi_{t s_T s_I})_{s_T = 1, ..., \bar{s}_T, s_I = 1, ..., \bar{s}_I} \in (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ where each $\chi_{t s_T s_I} \equiv (\chi_{t s_T s_I x})_{x = 1, ..., \bar{x}} \in \partial\Delta(\bar{x})$, we associate it with the x_t that satisfies

$$x_t(s_T, s_I, g) = x \text{ when } \sum_{x'=1}^{x-1} \chi_{t s_T s_I x'} \leq g < \sum_{x'=1}^{x} \chi_{t s_T s_I x'} \text{ for } x = 1, ..., \bar{x}, \tag{48}$$

and $x_t(s_T, s_I, 1) = \bar{x}$; also, we let $x_t(s_T, 0, g) = \infty$ which reflects the inventory-0-price-∞ association. We have effectively made each $\chi_{t s_T s_I x}$ the chance that choice x will be made by a type-s_T firm when in period t it is with inventory level $s_I \geq 1$. Now the search for equilibria $x^*_{[1\bar{t}]} \equiv (x^*_t | t = 1, ..., \bar{t}) \in (\mathcal{M}(S \times G), X))^{\bar{t}}$ can be reduced to that for the corresponding equilibria $\chi^*_{[1\bar{t}]} \equiv (\chi^*_t | t = 1, ..., \bar{t}) \in ((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$.

Given nonatomic game $\Gamma(\sigma_1)$, where $\sigma_1 \equiv (\sigma_{1 s_T s_I})_{s_T = 1, ..., \bar{s}_T, s_I = 0, 1, ..., \bar{s}_I} \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ is any initial pre-action environment that defines the game, analysis made in Appendix D tells us that we can construct a correspondence $\mathbf{Z}^D_D(\cdot|\sigma_1)$ on the space $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$. A fixed point $\chi^*_{[1\bar{t}]}$ satisfying $\chi^*_{[1\bar{t}]} \in \mathbf{Z}^D_D(\chi^*_{[1\bar{t}]}|\sigma_1)$ will provide a Markov equilibrium $x^*_{[1\bar{t}]} \in (\mathcal{M}(S \times G, X))^{\bar{t}}$ to the game. Meanwhile, existence of fixed points can be verified.

Proposition 5 *For any $\sigma_1 \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$, the correspondence $\mathbf{Z}^D_D(\cdot|\sigma_1)$ has fixed points.*

With discrete type, inventory, and price spaces S_T, S_I, and X, the $\tilde{s}_t$ defined at (47) automatically satisfies (S1). From (47), we know $\tilde{s}_{I,t}(s_T, s_I, x, \mu, i) \neq \tilde{s}_{I,t}(s_T, s_I, x, \mu', i)$ can occur only when $\tilde{\lambda}_{s_T x}(\mu|_{S_T \times X}) \leq i \leq \tilde{\lambda}_{s_T x}(\mu'|_{S_T \times X})$ or $\tilde{\lambda}_{s_T x}(\mu'|_{S_T \times X}) \leq i \leq \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X})$. Due to the finiteness of the spaces, (P1), (45), and the fact that $\bar{\iota}$ is uniform on $I \equiv [0, 1]$, this chance will go to zero at a speed independent of the s_T, s_I, and x present, as μ' converges to μ. So $\tilde{s}_t$ satisfies (S2) as well. With discrete type, inventory, and price spaces S_T, S_I, and X, the $\tilde{f}_t$ defined at (46) automatically satisfies (F1). By the finiteness of the spaces, (P1), and (45), it is clear that $\tilde{f}_t(s_T, s_I, x, \mu')$ will converge to $\tilde{f}_t(s_T, s_I, x, \mu)$ at a speed independent of the s_T, s_I, and x present, as μ' converges to μ. So $\tilde{f}_t$ satisfies (F2) as well.

The above indicate that the dynamic pricing game is compatible with the setting of Sections 3 to 6. For Theorem 2 to be applicable, we only need that the $x_t^* \in \mathcal{M}(S \times G, X)$ corresponding to each obtained equilibrium component $\chi_t^* \in (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ through (48) be a member of $\mathcal{K}(S, G, \bar{\gamma}, X)$. Due to (26), the latter boils down to that the prescribed price choice changes continuously in a $\bar{\gamma}$-measured probabilistic sense as the type or inventory level changes. But with discrete spaces S_T and S_I, this is automatically true.

Thus, when there are many of them, heterogeneous firms engaged in the aforementioned dynamic price competition do not have to fret about their inabilities to observe inventory levels of other firms. They can all adopt the same pricing policy, basically a map from their types, inventory levels, and individualized action-inducing shocks to price choices, as implied by an observation-blind equilibrium of the NG counterpart, and still expect to receive a total payoff that cannot on average be much improved by any one-time unilateral deviation. Note that the NG equilibrium may in turn be obtained from solving a fixed point problem, which in practice, can most likely be realized through an iterative procedure.

8.3 The Stationary Game

We now add production to the pricing game and study its stationary version. At the cost of $\bar{k}_{s_T}$, we suppose a type-s_T firm's inventory will be brought back to the full $\bar{s}_I$-level in one period's time when it has dropped down to 0. That is, the production lead time is one period and every firm exercises an (s^*, S^*) production policy with $s^* = 0$ and $S^* = \bar{s}_I$. The situation with $s^* \geq 1$ can be similarly handled. On top of the previous setup, we add a per-period discount factor $\bar{\phi} \in [0, 1)$.

For single-period payoff, we let $\tilde{f} : S_T \times S_I \times X \times \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x}) \longrightarrow \Re$ satisfy

$$\tilde{f}(s_T, s_I, x, \mu) = \mathbf{1}(s_I \geq 1) \cdot \bar{p}_x \cdot \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X}) - \bar{h}_{s_T} \cdot s_I - \bar{k}_{s_T} \cdot \mathbf{1}(s_I = 0). \tag{49}$$

The first two terms remain the same as those in (46). The extra term $\bar{k}_{s_T} \cdot \mathbf{1}(s_I = 0)$ reflects the total production cost needed to bring a type-s_T firm's inventory level back to $\bar{s}_I$ from 0. For state transition, we let $\tilde{s} \equiv (s_T, s_I) : S_T \times S_I \times X \times \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x}) \times I \longrightarrow S_T \times S_I$ satisfy

$$\begin{cases} \tilde{s}_T(s_T, s_I, x, \mu, i) = s_T, \\ \tilde{s}_I(s_T, s_I, x, \mu, i) = \mathbf{1}(s_I \geq 1) \cdot [s_I - \mathbf{1}(i \leq \tilde{\lambda}_{s_T x}(\mu|_{S_T \times X}))] + \mathbf{1}(s_I = 0) \cdot \bar{s}_I. \end{cases} \tag{50}$$

The T-portion of (50) is the same as that of (47). In the current I-portion, the first term has its counterpart in (47). The new term $\mathbf{1}(s_I = 0) \cdot \bar{s}_I$ reflects the locked-in production policy of bringing the inventory level back to full in one period's time.

For the stationary nonatomic version Γ of the dynamic pricing game, analysis made in Appendix E tells us that we can construct a correspondence $\mathbf{Z}^{DEF}_{DEF}$ on the space $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x}) \times \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$. A fixed point $(\chi^*, \xi^*, \sigma^*)$ satisfying $(\chi^*, \xi^*, \sigma^*) \in \mathbf{Z}^{DEF}_{DEF}(\chi^*, \xi^*, \sigma^*)$ will provide a Markov equilibrium $x^* \in \mathcal{M}(S \times G, X)$, all the while with σ^* serving as an associated pre-action environment. Again, existence of fixed points can be verified.

Proposition 6 *The correspondence $\mathbf{Z}^{DEF}_{DEF}$ has fixed points.*

Using similar arguments as those employed after Proposition 5 of Section 8.2, we can verify that (S1-s), (S2-s), (F1-s), and (F2-s) are all true. Also, the $x^* \in \mathcal{M}(S \times G, X)$ corresponding to each obtained equilibrium component $\chi^* \in (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ through (48) is a member of $\mathcal{K}(S, G, \bar{\gamma}, X)$. With Theorem 3 of Section 7 thus applicable, we can state the following. When production is allowed, we can learn from Proposition 6 about the existence of an NG equilibrium that prescribes for firms the dynamic pricing policy x^*. The latter would help the firms fare well in the long run under the environment σ^* which is consistent with all these firms adopting the policy x^*. When their number is finite but large enough, firms can still use the same pricing policy x^* to fare relatively well in the long run in an average sense that is defined through the joint type-inventory distribution σ^*.

9 Discussion

Using this paper's language and notation, we offer a comparison with the most relevant papers. Within the discrete-time framework while without considering atomic players or

players' entries and exits, we have arguably worked with the most general setup. We also touch on the existence of NG equilibria.

9.1 Comparison with Most Relevant Literature

Both Weintraub, Benkard, and van Roy [35] and Weintraub et al. [33] treated competing firms on a common market as players. They allowed for entry and exit of firms, and accounted for the effect of firm density c per unit market size. Roughly speaking, their regular payoff is in the form of $\tilde{f}^0(s, c \cdot \mu|_S) - \tilde{f}^1(x)$, where $\mu|_S$ stands for the marginal state distribution derivable from the joint state-action distribution μ. Also, firms' state transitions are governed by a certain $\tilde{s}^0(s, x, i)$ that is independent of the environment μ. In finite models, $\tilde{f}^0$ is replaced by some $f^{0\prime}(s, \mu|_S, n, m)$ where n is the number of firms and m the market size. In the limiting regime, it was certainly assumed that

$$\lim_{m \to +\infty} \tilde{f}^{0\prime}(s, \mu|_S, n(m), m) = \tilde{f}^0(s, c \cdot \mu|_S), \tag{51}$$

when $\lim_{m \to +\infty} n(m)/m = c$.

Weintraub, Benkard, and van Roy [35] found a stationary policy of the form $x(s)$ to suffice for the NG limit. It was considered oblivious because of firms' abilities to ignore the industry state $c \cdot \mu|_S$. When there are few dominant firms in it, an NG equilibrium was shown to work increasingly well for larger finite models. This is close in spirit to our Theorem 3. In the mean time, Weintraub et al. [33] arrived something akin to our Theorem 2. We note that $\tilde{s}^0$'s independence of μ helped greatly with their derivations. While free from the task of dealing with entry, exit, or impacts of market size and number of firms, we have allowed players' state transitions to be profoundly impacted by the environment that their collective states and actions fabricate. Namely, our $\tilde{s}_t$ can depend on μ in virtually arbitrary fashions. Meanwhile, according to (47) and (50), such dependence is integral to the pricing game.

Huang, Malhame, and Caines [16] dealt with continuous-time games with the state space S equal to the real line $\Re$. These games' discrete-time counterparts can be obtained by replacing their Brownian motions with symmetric random walks. In particular, we can let the transition-inducing shock space I be $\{-1, +1\}$ and the probability $\bar{\iota}$ be half on -1 and half on $+1$. When thus cast, the earlier work's state transition can be understood as

$$\tilde{s}_t(s, x, \mu, i) = \int_{\Re} \tilde{s}^0(s, x, s') \cdot \mu|_S(ds') + \bar{s}^1 \cdot i, \tag{52}$$

where $\tilde{s}^0$ is a function from $\Re \times X \times \Re$ to $\Re$ and $\bar{s}^1$ is a constant. So there, only the state-distribution portion of the joint state-action distribution μ of other firms affect the current

firm's state transition; its impact is also felt in an average sense; moreover, the effect of the random shock is additive. This certainly did not cover the case of (47) or (50).

Their one-period payoff function can be understood as

$$\tilde{f}_t(s, x, \mu) = \int_{\Re} \tilde{f}^0(s, x, s') \cdot \mu|_S(ds'), \tag{53}$$

where $\tilde{f}^0$ is a function from $\Re \times X \times \Re$ to $\Re$. Therefore, payoffs depend on in-action environments also in a special fashion incompatible with (46) and (49) used in the dynamic pricing game. Artificial randomization in decision making turns out to be unnecessary—NG equilibria can be found in the form of $x_t(s)$ rather than the more general $x_t(s, g)$.

We, on the other hand, believe that allowing other players' actions to play a role in both state transitions and one-period payoffs can greatly enhance the relevant models' applicabilities. In the competitive pricing situation, for instance, the demand level experienced by a firm is perturbable by prices charged by other firms. It in turn influences not only the firm's present profitability but also its future inventory levels.

Like this paper, Yang [38] also showed that pre-action environments of large finite games will converge to those of their NG counterparts, and therefore NG equilibria can be put to good use in large finite games in both transient and stationary cases. Yet, differences on space requirements set the two works apart. Derivations in Yang [38] rely on discrete metric state and action spaces S and X. In the current case, S only has to be complete separable and X separable. The restriction we place in turn is the identifiability of shocks that drive both random decision makings and stochastic state transitions. As shown around (5) and (8), these are, for instance, achievable when S and X are closed subsets of the real line.

9.2 Existence of NG Equilibria

NG equilibria's existence is not our focal point. Their usefulness in finite-player situations is. Briefly, though, we show that NG equilibria are guaranteed to exist when state and action spaces are discrete. We consider a metric space discrete when it has a countable number of elements and the minimum distance between any two elements is strictly positive. Suppose that both the state space S and action space X are discrete subsets of the real line $\Re$, both shock spaces G and I are the interval $[0, 1]$, and both shock distributions $\bar{\gamma}$ and $\bar{\iota}$ are uniform.

Given pre-action environment $\sigma_t \in \mathcal{P}(S)$ and decision rule $x_t(s, g)$, we have defined the induced in-action environment through (9) and (10). Given the same σ_t and its decision

rule $\chi_t(s)$, Yang [38] gave the corresponding definition of $\sigma \otimes \chi_t$ in the sense that for any $S' \in \mathcal{B}(S)$ and $X' \in \mathcal{B}(X)$,

$$(\sigma \otimes \chi_t)(S' \times X') = \int_{S'} \sigma_t(ds) \cdot \chi_t(s|X'). \tag{54}$$

Proposition 7 *The in-action environments considered in the two papers are the same.*

Technical details of this section are relegated to Appendix F. Given decision rule $x_t(s, g)$ and implicitly the state transition rule $\tilde{s}_t(s, x, \mu, i)$, we have defined the pre-action environment transition operator $T_t(x_t)$ through (12) and (13). Given its decision rule $\chi_t(s)$ and implicit state transition rule $\tilde{g}_t(s, x, \mu)$, Yang [38] defined a corresponding operator $\hat{T}_t(\chi_t)$. For any $\sigma_t \in \mathcal{P}(S)$ and $S' \in \mathcal{B}(S)$,

$$[\hat{T}_t(\chi_t) \circ \sigma_t](S') = \int_S \sigma_t(ds) \cdot \int_X \chi_t(s|dx) \cdot \tilde{g}_t(s, x, \sigma_t \otimes \chi_t|S'). \tag{55}$$

Proposition 8 *The two pre-action environment transition operators are equivalent.*

In the current discrete-space setting, with the equivalences provided by Propositions 7 and 8, the equilibrium notions defined in the two papers are almost identical, with the exception being that our current definitions are slightly weaker due to an average over the current players' state. For instance, one may compare the transient NG criterion (19) with Yang's [38] Section 6 and the stationary NG criterion (34) with his Section 8.

Theorem 5 of Yang [38] has shown the existence of transient NG equilibria $\chi^* = (\chi_t^* | t = 1, ..., \bar{t})$ along with associated pre-action environments $\sigma_2^* = \hat{T}_1(\chi_1^*) \circ \sigma_1$, $\sigma_3^* = \hat{T}_2(\chi_2^*) \circ \sigma_2^*$, ..., $\sigma_{\bar{t}+1}^* = \hat{T}_{\bar{t}}(\chi_{\bar{t}}^*) \circ \sigma_{\bar{t}}^*$. Applying (8) on every χ_t^* for $t = 1, ..., \bar{t}$, we can obtain an equilibrium $x^* = (x_t^* | t = 1, ..., \bar{t})$ for the current situation. Similarly, Theorem 6 of the existing work has shown the existence of a stationary NG equilibria χ^* along with associated pre-action environments σ^*. We can obtain an equilibrium x^* for the current situation by applying the time-invariant version of (8) on χ^*.

Transient and stationary policies obtained in the above fashion are monotone and hence measurable in g. As S is discrete, they are all members of $\mathcal{K}(S, G, \bar{\gamma}, X)$ and hence $\mathcal{M}(S \times G, X)$. For more general spaces, we may deal with equilibrium existence on a case-by-case basis, with Jovanovic and Rosenthal [18] and Bergin and Bernhardt [6] offering inspirations.

10 Concluding Remarks

Using the weak LLN for empirical distributions, we have established links between multi-period Markovian games and their NG counterparts. Our focus is the case where state and action spaces are general metric spaces, and there are independently generated shocks serving as random drivers for decision making and state evolution. In essence, the evolution of player-state distributions in large finite games, though random, resembles in probability the deterministic pathway taken by their NG brethren. This allows NG equilibria to be well adapted to large finite games. Our results have been successfully brought to bear on dynamic pricing situations with their own practical significances.

It is essential in our derivations that both action- and transition-inducing shocks be independent of one another. When there is a global shock commonly experienced by all players, then the dependences of payoffs and transition functions on this shock will translate into inter-dependences among different players. It will not take too much to adapt our approach to this situation. There, the environment pathway of even an NG will not be deterministic. Rather, it will be dependent on the global shocks just experienced.

Still, many dynamic competitive situations not yet covered by existing studies like Huang, Malhame, and Caines [16] are better described by continuous-time models. These will require vastly different techniques to probe. For one thing, the mathematical induction approach we have taken to deal with multiple periods would not seem to go well with a discrete-time approximation of a continuous-time model. In the latter model, even to identify the environment induced by all players adopting a common policy might involve solving a fixed point problem. Therefore, serious challenges will have to be overcome.

Appendices

A Technical Developments in Section 6.1

Given metric space A, the Prohorov metric ρ_A is such that, for any distributions $\pi, \pi' \in \mathcal{P}(A)$,

$$\rho_A(\pi, \pi') = \inf(\epsilon > 0 \mid \pi'((A')^\epsilon) + \epsilon \geq \pi(A'), \quad \text{for all } A' \in \mathcal{B}(A)), \tag{A.1}$$

where

$$(A')^\epsilon = \{a \in A \mid d_A(a, a') < \epsilon \text{ for some } a' \in A'\}. \tag{A.2}$$

The metric ρ_A is known to generate the weak topology for $\mathcal{P}(A)$.

Proof of Proposition 1: We prove (I) first. Let $\epsilon > 0$ be given. By the hypothesis, we know there is $\delta > 0$, so that for any $a, a' \in A$ with $d_A(a, a') < \delta$, for the set

$$B' = \{b \in B | d_C(y(a, b), y(a', b)) < \epsilon\} \in \mathcal{B}(B), \tag{A.3}$$

we have

$$\pi_B(B') > 1 - \epsilon. \tag{A.4}$$

Now consider any $C' \in \mathcal{B}(C)$. The key observation from (A.3) is that

$$(y(a, \cdot))^{-1}(C') \cap B' \subseteq (y(a', \cdot))^{-1}((C')^{\epsilon}). \tag{A.5}$$

Then,

$$\begin{aligned} \theta(a|C') &= \pi_B((y(a, \cdot))^{-1}(C') \cap B') + \pi_B((y(a, \cdot))^{-1}(C') \cap (B \setminus B')) \\ &\leq \pi_B((y(a', \cdot))^{-1}((C')^{\epsilon})) + \pi_B(B \setminus B') < \theta(a'|(C')^{\epsilon}) + \epsilon, \end{aligned} \tag{A.6}$$

where the only equality is an identity, the first inequality is due to (A.5), and the second inequality is due to (A.4). So by (A.1), we have $\rho_C(\theta(a), \theta(a')) < \epsilon$.

We prove (II) next. By y's definition, we have

$$b \in [0, \theta(a|(-\infty, c])] \quad \text{if and only if} \quad y(a, b) \leq c. \tag{A.7}$$

Thus,

$$\pi_B(\{b \in B | y(a, b) \leq c\}) = \pi_B([0, \theta(a|(-\infty, c])]) = \theta(a|(-\infty, c]), \tag{A.8}$$

where the first equality is from (A.7) and the second one is from the uniform nature of π_B. So indeed $\pi_B \cdot (y(a, \cdot))^{-1}$ is $\theta(a)$.

Let sequence a_n converge to $a \in A$. By $\theta(\cdot)$'s continuity at a, we know that $\pi_B \cdot (y(a_n, \cdot))^{-1}$ converges to $\pi_B \cdot (y(a, \cdot))^{-1}$ under the ρ_C-metric. Since B is $[0, 1]$, π_B is uniform, and C is complete, we know from Skorohod's representation theorem, that there exist functions $z_n \in \mathcal{M}(B, C)$ and $z \in \mathcal{M}(B, C)$ with $\pi_B \cdot z_n^{-1} = \pi_B \cdot (y(a, \cdot))^{-1}$ and $z^{-1} = \pi_B \cdot (y(a, \cdot))^{-1}$ such that z_n converges to z almost surely.

According to Jouini, Schachermayer, and Touzi [17] (Lemma A.4), there is a one-to-one measure-preserving mapping $e : B \longrightarrow B$ so that $y(a, b) = z(e(b))$. Note $z_n(e(\cdot))$ converges to $z(e(\cdot)) = y(a, \cdot)$ almost surely as well. Without loss of generality, we suppose e is the identity mapping on B. This entails that, $z_n(\cdot)$, which has the same distribution as $y(a_n, \cdot)$, converges to the monotone $z(\cdot) = y(a, \cdot)$.

Let $\epsilon > 0$ be given. Because $y(a, \cdot)$ is continuous on the compact $B = [0, 1]$ at an a-independent rate, there exists $\delta \in (0, \epsilon]$, so that for $b, b' \in B$ with $\mid b - b' \mid < \delta$ and any $n \in N$, one has

$$\mid y(a_n, b) - y(a_n, b') \mid < \frac{\epsilon}{2}. \tag{A.9}$$

Due to z_n's almost sure convergence to z, we also know that, for n large enough,

$$\pi_B(B'_n) > 1 - \delta \geq 1 - \epsilon, \tag{A.10}$$

where

$$B'_n = \{b \in B | \mid z_n(b) - z(b) \mid < \frac{\epsilon}{2}\}. \tag{A.11}$$

There is a measure-preserving map f_n on B'_n so that $z_n(f_n(\cdot))$ is monotone on B'_n. Since $z(\cdot) = y(a, \cdot)$ is monotone, one key observation from (A.11) is that

$$\mid z_n(f_n(b)) - z(b) \mid < \frac{\epsilon}{2}, \qquad \forall b \in B'_n. \tag{A.12}$$

Since $z_n(\cdot)$ has the same distribution as $y(a_n, \cdot)$, there is a measure-preserving map g_n on B so that $z_n(g_n(\cdot)) = y(a_n, \cdot)$. Due to $y(a_n, \cdot)$'s monotonicity, $z_n(g_n(\cdot))$ is monotone on the entire B.

By the monotonicity of $z_n(f_n(\cdot))$ on B'_n, that of $z_n(g_n(\cdot))$ on B, and (A.10), we reach another key observation that

$$\mid f_n(b) - g_n(b) \mid < \delta, \qquad \forall b \in B'_n. \tag{A.13}$$

Now for $b \in B'_n$,

$$\mid y(a_n, b) - y(a, b) \mid = \mid z_n(g_n(b)) - z(b) \mid \leq \mid z_n(g_n(b)) - z_n(f_n(b)) \mid + \mid z_n(f_n(b)) - z(b) \mid . \tag{A.14}$$

By (A.9) and (A.13), the first term on the right-hand side is below $\epsilon/2$; whereas, by (A.12), the second term on the right-hand side is below $\epsilon/2$ as well. It is now clear that $y(a_n, \cdot)$ converges to $y(a, \cdot)$ in probability. ∎

According to Parthasarathy [26] (Theorem II.7.1), the strong LLN applies to the empirical distribution under the weak topology, and hence under the Prohorov metric. In the following, we state its weak version.

Lemma 1 *Given separable metric spaces A and B, suppose distribution $\pi_A \in \mathcal{P}(A)$ and measurable mapping $y \in \mathcal{M}(A, B)$. Then, for any $\epsilon > 0$, as long as n is large enough,*

$$(\pi_A)^n(\{a = (a_1, ..., a_n) \in A^n \mid \rho_B(\varepsilon_a \cdot y^{-1}, \pi \cdot y^{-1}) < \epsilon\}) > 1 - \epsilon.$$

For separable metric space A, point $a \in A$, and the $(n-1)$-point empirical distribution space $\pi \in \mathcal{P}_{n-1}(A)$, we use $(a,\pi)_n$ to represent the member of $\mathcal{P}_n(A)$ that has an additional $1/n$ weight on the point a, but with probability masses in π being reduced to $(n-1)/n$ times of their original values. For $a \in A^n$ and $m = 1, ..., n$, we have $(a_m, \varepsilon_{a_{-m}})_n = \varepsilon_a$. Concerning the Prohorov metric, we have also a simple but useful observation.

Lemma 2 *Let A be a separable metric space. Then, for any $n = 2, 3, ...$, $a \in A$, and $\pi \in \mathcal{P}_{n-1}(A)$,*

$$\rho_A((a,\pi)_n, \pi) \leq \frac{1}{n}.$$

Proof: Let $A' \in \mathcal{B}(A)$ be chosen. If $a \notin A'$, then

$$(a,\pi)_n(A') \leq \pi(A') \leq (a,\pi)_n(A') + \frac{1}{n}; \tag{A.15}$$

if $a \in A'$, then

$$(a,\pi)_n(A') - \frac{1}{n} \leq \pi(A') \leq (a,\pi)_n(A'). \tag{A.16}$$

Hence, it is always true that

$$\mid (a,\pi)_n(A') - \pi(A') \mid \leq \frac{1}{n}. \tag{A.17}$$

In view of (A.1) and (A.2), we have

$$\rho_A((a,\pi)_n, \pi) \leq \frac{1}{n}. \tag{A.18}$$

We have thus completed the proof. ∎

The following result is important for showing the near-trajectory evolution of aggregate environments in large multi-period games.

Lemma 3 *Given separable metric space A and complete separable metric spaces B and C, suppose $y_n \in \mathcal{M}(A^n, B^n)$ for every $n \in N$, $\pi_A \in \mathcal{P}(A)$, $\pi_B \in \mathcal{P}(B)$, and $\pi_C \in \mathcal{P}(C)$. If*

$$(\pi_A)^n(\{a \in A^n \mid \rho_B(\varepsilon_{y_n(a)}, \pi_B) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough, then

$$(\pi_A \times \pi_C)^n(\{(a,c) \in (A \times C)^n \mid \rho_{B\times C}(\varepsilon_{y_n(a),c}, \pi_B \times \pi_C) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough.

Proof: Suppose sequence $\{\pi'_{B1}, \pi'_{B2}, ...\}$ weakly converges to the given probability measure π_B, and sequence $\{\pi'_{C1}, \pi'_{C2}, ...\}$ weakly converges to the given probability measure π_C. We are to show that the sequence $\{\pi'_{B1} \times \pi'_{C1}, \pi'_{B2} \times \pi'_{C2}, ...\}$ weakly converges to $\pi_B \times \pi_C$.

Let $F(B)$ denote the family of uniformly continuous real-valued functions on B with bounded support. Let $F(C)$ be similarly defined for C. We certainly have

$$\begin{cases} \lim_{k\to+\infty} \int_B f(b) \cdot \pi'_{Bk}(db) = \int_B f(b) \cdot \pi_B(db), & \forall f \in F(B), \\ \lim_{k\to+\infty} \int_C f(c) \cdot \pi'_{Ck}(dc) = \int_C f(c) \cdot \pi_C(dc), & \forall f \in F(C). \end{cases} \tag{A.19}$$

Define F so that

$$\begin{aligned} F = \{f \mid f(b,c) = f_B(b) \cdot f_C(c) \ \text{ for any } (b,c) \in B \times C, \\ \text{where } f_B \in F(B) \cup \{\mathbf{1}\} \text{ and } f_C \in F(C) \cup \{\mathbf{1}\}\}, \end{aligned} \tag{A.20}$$

where $\mathbf{1}$ stands for the function whose value is 1 everywhere. By (A.19) and (A.20),

$$\lim_{k\to+\infty} \int_{B\times C} f(b,c) \cdot (\pi'_{Bk} \times \pi'_{Ck})(d(b,c)) = \int_{B\times C} f(b,c) \cdot (\pi_B \times \pi_C)(d(b,c)). \tag{A.21}$$

According to Ethier and Kurtz [9] (Proposition III.4.4), $F(B)$ and $F(C)$ happen to be $\mathcal{P}(B)$ and $\mathcal{P}(C)$'s convergence determining families, respectively. As B and C are complete, Ethier and Kurtz ([9], Proposition III.4.6, whose proof involves Prohorov's Theorem, i.e., the equivalence between tightness and relative compactness of a collection of probability measures defined for complete separable metric spaces) further states that F as defined through (A.20) is convergence determining for $\mathcal{P}(B \times C)$. Therefore, we have the desired weak convergence by (A.21).

Let $\epsilon > 0$ be given. In view of the above product-measure convergence and the equivalence between the weak topology and that induced by the Prohorov metric, there must be $\delta_B > 0$ and $\delta_C > 0$, such that $\rho_B(\pi'_B, \pi_B) < \delta_B$ and $\rho_C(\pi'_C, \pi_C) < \delta_C$ will imply

$$(\rho_B \times \rho_C)(\pi'_B \times \pi'_C, \pi_B \times \pi_C) < \epsilon. \tag{A.22}$$

By (A.1) and the given hypothesis, there is $\bar{n}^1 \in N$, so that for $n = \bar{n}^1, \bar{n}^1 + 1, ...,$

$$(\pi_A)^n(\tilde{A}_n) > 1 - \frac{\epsilon}{2}, \tag{A.23}$$

where $\tilde{A}_n$ contains all $a \in A^n$ such that

$$\rho_B(\varepsilon_{y_n(a)}, \pi_B) < \delta_B. \tag{A.24}$$

By (A.1) and Lemma 1, on the other hand, there is $\bar{n}^2 \in N$, so that for $n = \bar{n}^2, \bar{n}^2 + 1, ...,$

$$(\pi_C)^n(\tilde{C}_n) > 1 - \frac{\epsilon}{2}, \tag{A.25}$$

where $\tilde{C}_n$ contains all $c \in C^n$ such that

$$\rho_C(\varepsilon_c, \pi_C) < \delta_C. \tag{A.26}$$

For any $n = \bar{n}^1 \vee \bar{n}^2, \bar{n}^1 \vee \bar{n}^2 + 1, ...,$ let (a, c) be an arbitrary member of $\tilde{A}_n \times \tilde{C}_n$. We have from (A.22), (A.24), and (A.26) that,

$$(\rho_B \times \rho_C)(\varepsilon_{y_n(a),c}, \pi_B \times \pi_C) < \epsilon. \tag{A.27}$$

Noting the facilitating (a, c) is but an arbitrary member of $\tilde{A}_n \times \tilde{C}_n$, we see that

$$\begin{aligned}(\pi_A \times \pi_C)^n(\{(a,c) \in (A \times C)^n \mid \rho_{B\times C}(\varepsilon_{y_n(a),c}, \pi_B \times \pi_C) < \epsilon\}) \\ \geq (\pi_A)^n(\tilde{A}_n) \times (\pi_C)^n(\tilde{C}_n),\end{aligned} \tag{A.28}$$

which by (A.23) and (A.25), is greater than $1 - \epsilon$. ∎

Because the equivalence between tightness and relative compactness of a collection of probability measures is indirectly related to the proof of Lemma 3, we require B and C to be complete separable metric spaces.

Lemma 4 *Given separable metric spaces A, B, C, and D, as well as distributions $\pi_A \in \mathcal{P}(A)$, $\pi_B \in \mathcal{P}(B)$, and $\pi_C \in \mathcal{P}(C)$, suppose $y_n \in \mathcal{M}(A^n, B^n)$ for every $n \in N$ and $z \in \mathcal{K}(B, C, \pi_C, D)$. If*

$$(\pi_A \times \pi_C)^n(\{a \in A^n, c \in C^n \mid \rho_{B\times C}(\varepsilon_{y_n(a),c}, \pi_B \times \pi_C) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough, then

$$(\pi_A \times \pi_C)^n(\{a \in A^n, c \in C^n \mid \rho_D(\varepsilon_{y_n(a),c} \cdot z^{-1}, (\pi_B \times \pi_C) \cdot z^{-1}) < \epsilon\}) > 1 - \epsilon,$$

for any $\epsilon > 0$ and any n large enough.

Proof: Let $\epsilon > 0$ be given. Since $z \in \mathcal{K}(B, C, \pi_C, D)$, there exist $C' \in \mathcal{B}(C)$ satisfying

$$\pi_C(C') > 1 - \frac{\epsilon}{2}, \tag{A.29}$$

as well as

$$\delta \in (0, \epsilon/2], \tag{A.30}$$

such that for any $b, b' \in B$ satisfying $d_B(b, b') < \delta$ and any $c \in C'$,

$$d_D(z(b, c), z(b', c)) < \epsilon. \tag{A.31}$$

For any subset D' in $\mathcal{B}(D)$, we therefore have

$$(z^{-1}(D'))^\delta \cap (B \times C') \subseteq z^{-1}((D')^\epsilon). \tag{A.32}$$

This leads to $(z^{-1}(D'))^\delta \setminus (B \times (C \setminus C')) \subseteq z^{-1}((D')^\epsilon)$, and hence due to (A.29),

$$(\pi_B \times \pi_C)(z^{-1}((D')^\epsilon)) \geq (\pi_B \times \pi_C)((z^{-1}(D'))^\delta) - \frac{\epsilon}{2}. \tag{A.33}$$

On the other hand, by the hypothesis, we know for n large enough,

$$(\pi_A \times \pi_C)^n(E'_n) > 1 - \delta, \tag{A.34}$$

where

$$E'_n = \{a \in A^n, c \in C^n \mid \rho_{B\times C}(\varepsilon_{y_n(a),c}, \pi_B \times \pi_C) < \delta\} \in \mathcal{B}^n(A \times C). \tag{A.35}$$

By (A.35), for any $(a, b) \in E'_n$ and $F' \in \mathcal{B}(B \times C)$,

$$(\pi_B \times \pi_C)((F')^\delta) \geq \varepsilon_{y_n(a),c}(F') - \delta. \tag{A.36}$$

Combining the above, we have, for any $(a, c) \in E'_n$ and $D' \in \mathcal{B}(D)$,

$$\begin{aligned} [(\pi_B \times \pi_C) \cdot z^{-1}]((D')^\epsilon) &= (\pi_B \times \pi_C)(z^{-1}((D')^\epsilon)) \\ &\geq (\pi_B \times \pi_C)((z^{-1}(D'))^\delta) - \epsilon/2 \geq \varepsilon_{y_n(a),c}(z^{-1}(D')) - \delta - \epsilon/2 \\ &\geq \varepsilon_{y_n(a),c}(z^{-1}(D')) - \epsilon = (\varepsilon_{y_n(a),c} \cdot z^{-1})(D') - \epsilon. \end{aligned} \tag{A.37}$$

where the first inequality is due to (A.33), the second inequality is due to (A.36), and the third inequality is due to (A.30). That is, we have

$$\rho_D(\varepsilon_{y_n(a),c} \cdot z^{-1}, (\pi_B \times \pi_C) \cdot z^{-1}) \leq \epsilon, \qquad \forall (a, c) \in E'_n. \tag{A.38}$$

In view of (A.30) and (A.34), we have the desired result. ∎

We can now prove Proposition 2 and then Theorem 1.

Proof of Proposition 2: Let $t = 1, ..., \bar{t} - 1$ and $x \in \mathcal{K}(S, G, \bar{\gamma}, X)$ be given. Define map $z \in \mathcal{M}(S \times G \times I, S)$, so that

$$z(s, g, i) = \tilde{s}_t(s, x(s, g), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, i), \qquad \forall s \in S, g \in G, i \in I. \tag{A.39}$$

In view of (12) and (A.39), we have, for any $S' \in \mathcal{B}(S)$,

$$\begin{aligned}[T_t(x) \circ \sigma](S') &= \textstyle\int_S \int_G \int_I \mathbf{1}(z(s, g, i) \in S') \cdot \bar{\iota}(di) \cdot \bar{\gamma}(dg) \cdot \sigma(ds) \\ &= (\sigma \times \bar{\gamma} \times \bar{\iota})(\{(s, g, i) \in S \times G \times I \mid z(s, g, i) \in S'\}) = (\sigma \times \bar{\gamma} \times \bar{\iota})(z^{-1}(S')).\end{aligned} \tag{A.40}$$

For $n \in N$, $g = (g_1, ..., g_n) \in G^n$, and $i = (i_1, ..., i_n) \in I^n$, also define operator $T'_n(g, i)$ on $\mathcal{P}_n(S)$ so that $T'_n(g, i) \circ 1_S = \varepsilon_{s'}$, where for $m = 1, 2, ..., n$,

$$s'_m = z(s_m, g_m, i_m) = \tilde{s}_t(s_m, x(s_m, g_m), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, i_m). \tag{A.41}$$

It is worth noting that (A.41) is different from the earlier (21). In view of (A.39) and (A.41), we have, for $S' \in \mathcal{B}(S)$,

$$[T'_n(g, i) \circ 1_S](S') = \frac{1}{n} \cdot \sum_{m=1}^{n} \mathbf{1}(z(s_m, g_m, i_m) \in S') = \varepsilon_{((s_1, g_1, i_1), ..., (s_n, g_n, i_n))}(z^{-1}(S')). \tag{A.42}$$

Combining (A.40) and (A.42), we arrive to a key observation that

$$T_t(x) \circ \sigma = (\sigma \times \bar{\gamma} \times \bar{\iota}) \cdot z^{-1}, \qquad \text{while} \quad T'_n(g, i) \circ 1_S = \varepsilon_{s,g,i} \cdot z^{-1}. \tag{A.43}$$

In the rest of the proof, we first show the asymptotic closeness between $T_t(x) \circ \sigma$ and $T'_n(g, i) \circ \varepsilon_{s_n(a)}$, and then that between the latter and $T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}$.

First, due to the hypothesis on the convergence of $\varepsilon_{s_n(a)}$ to σ, the completeness of the spaces S, G, and I and hence also the completeness of $G \times I$, as well as Lemma 3,

$$(\pi \times \bar{\gamma} \times \bar{\iota})^n(\{(a, g, i) \in (A \times G \times I)^n \mid \rho_{S \times G \times I}(\varepsilon_{s_n(a), g, i}, \sigma \times \bar{\gamma} \times \bar{\iota}) < \epsilon'\}) > 1 - \epsilon', \tag{A.44}$$

for any $\epsilon' > 0$ and any n large enough. By (S1) and the fact that $x \in \mathcal{K}(S, G, \bar{\gamma}, X)$, we may see that z as defined through (A.39) is a member of $\mathcal{K}(S, G \times I, \bar{\gamma} \times \bar{\iota}, S)$. By Lemma 4, this fact along with (A.44) will lead to

$$(\pi \times \bar{\gamma} \times \bar{\iota})^n(\{(a, g, i) \in (A \times G \times I)^n \mid \rho_S(\varepsilon_{s_n(a), g, i} \cdot z^{-1}, (\sigma \times \bar{\gamma} \times \bar{\iota}) \cdot z^{-1}) < \epsilon'\}) > 1 - \epsilon', \tag{A.45}$$

for any $\epsilon' > 0$ and any n large enough. By (A.43), this is equivalent to that, given $\epsilon > 0$, there exists $\bar{n}^1 \in N$ so that for any $n = \bar{n}^1, \bar{n}^1 + 1, ...,$

$$(\pi \times \bar{\gamma} \times \bar{\iota})^n(\tilde{A}_n(\epsilon)) > 1 - \frac{\epsilon}{2}, \tag{A.46}$$

where

$$\tilde{A}_n(\epsilon) = \{(a,g,i) \in (A \times G \times I)^n \mid \rho_S(T_t(x) \circ \sigma, T'_n(g,i) \circ \varepsilon_{s_n(a)}) < \frac{\epsilon}{2}\} \in \mathcal{B}^n(A \times G \times I). \quad \text{(A.47)}$$

Next, note that the only difference between $T_{nt}(x,g,i) \circ \varepsilon_{s_n(a)}$ and $T'_n(g,i) \circ \varepsilon_{s_n(a)}$ lies in that $\varepsilon_{s_{n,-m}(a),g_{-m}}$ is used in the former as in (21) whereas $\sigma \times \bar{\gamma}$ is used in the latter as in (A.41). Here, $s_{n,-m}(a)$ refers to the vector $(s_{n1}(a), ..., s_{n,m-1}(a), s_{n,m+1}(a), ..., s_{nn}(a))$. By (S2), there is $\delta \in (0, \epsilon/4]$ and $I' \in \mathcal{B}(I)$ with

$$\bar{\iota}(I') > 1 - \frac{\epsilon}{4}, \quad \text{(A.48)}$$

so that for any $(s,g,i) \in S \times G \times I'$ and any $\mu' \in \mathcal{P}(S \times X)$ satisfying $\rho_{S \times X}((\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, \mu') < \delta$,

$$d_S(\tilde{s}_t(s, x(s,g), (\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, i), \tilde{s}_t(s, x(s,g), \mu', i)) < \frac{\epsilon}{2}. \quad \text{(A.49)}$$

For each $n \in N$, define I'_n so that

$$I'_n = \{i = (i_1, ..., i_n) \in I^n \mid \text{more than } (1 - \frac{\epsilon}{2}) \cdot n \text{ components come from } I'\}. \quad \text{(A.50)}$$

Also important is that by (A.49) and (A.50), for any $S' \in \mathcal{B}(S)$ and $i = (i_1, ..., i_n) \in I'_n$,

$$[T_{nt}(x,g,i) \circ \varepsilon_{s_n(a)}]((S')^{\epsilon/2}) + \frac{\epsilon}{2} \geq [T'_n(g,i) \circ \varepsilon_{s_n(a)}](S'), \quad \text{(A.51)}$$

whenever

$$\rho_{S \times X}((\sigma \times \bar{\gamma}) \cdot (1_S, x)^{-1}, \varepsilon_{s_{n,-m}(a),g_{-m}} \cdot (1_S, x)^{-1}) < \delta. \quad \text{(A.52)}$$

It can be shown that I'_n will occupy a big chunk of I^n as measured by $\bar{\tau}^n$ when n is large. Define map q from I to $\{0,1\}$ so that $q(i) = 1$ or 0 depending on whether or not $i \in I'$. By (A.48), $\bar{\tau} \cdot q^{-1}$ is a Bernoulli distribution with $(\bar{\tau} \cdot q^{-1})(\{1\}) > 1 - \epsilon/4$. So by (A.50), I'_n contains all $i = (i_1, ..., i_n) \in I^n$ that satisfy

$$\rho_{\{0,1\}}(\varepsilon_i \cdot q^{-1}, \bar{\tau} \cdot q^{-1}) < \frac{\epsilon}{4}. \quad \text{(A.53)}$$

Therefore, by Lemma 1, there exits $\bar{n}^2 \in N$, so that for $n = \bar{n}^2, \bar{n}^2 + 1, ...,$

$$\bar{\iota}^n(I'_n) > 1 - \frac{\epsilon}{4}. \quad \text{(A.54)}$$

We can also demonstrate that (A.52) will be highly likely when n is large. By Lemma 3 and the hypothesis on the convergence of $\varepsilon_{s_n(a)}$ to σ, we know $\varepsilon_{s_n(a),g}$ will converge to $\sigma \times \bar{\gamma}$

in probability. Due to Lemma 2, this conclusion applies to the sequence $\varepsilon_{s_{n,-m}(a),g_{-m}}$ as well. The fact that $x \in \mathcal{K}(S, G, \bar{\gamma}, X)$ certainly leads to $(1_S, x) \in \mathcal{K}(S, G, \bar{\gamma}, S \times X)$. So by Lemma 4, there is $\bar{n}^3 \in N$, so that for $n = \bar{n}^3, \bar{n}^3 + 1, ...,$

$$(\pi^n \times \bar{\gamma}^n)(\tilde{B}_n(\delta)) > 1 - \frac{\epsilon}{4}, \tag{A.55}$$

where

$$\tilde{B}_n(\delta) = \{(a, g) \in A^n \times G^n \mid \text{(A.52) is true}\} \in \mathcal{B}^n(A \times G). \tag{A.56}$$

Consider arbitrary $n = \bar{n}^1 \vee \bar{n}^2 \vee \bar{n}^3, \bar{n}^1 \vee \bar{n}^2 \vee \bar{n}^3 + 1, ...,$ $(a, g, i) \in \tilde{A}_n(\epsilon) \cap (\tilde{B}_n(\delta) \times I'_n)$, and $S' \in \mathcal{B}(S)$. By (A.1) and (A.47), we see that

$$[T'_n(g, i) \circ \varepsilon_{s_n(a)}]((S')^{\epsilon/2}) + \frac{\epsilon}{2} \geq [T_t(x) \circ \sigma](S'). \tag{A.57}$$

Combining this with (A.51), (A.52), and (A.56), we obtain

$$[T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}]((S')^{\epsilon}) + \epsilon \geq [T'_n(g, i) \circ \varepsilon_{s_n(a)}]((S')^{\epsilon/2}) + \frac{\epsilon}{2} \geq [T_t(x) \circ \sigma](S'). \tag{A.58}$$

According to (A.1), this means

$$\rho_S(T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}, T_t(x) \circ \sigma) \leq \epsilon. \tag{A.59}$$

Therefore, for $n \geq \bar{n}^1 \vee \bar{n}^2 \vee \bar{n}^3$,

$$\begin{aligned}(\pi \times \bar{\gamma} \times \bar{\iota})^n(\{(a, g, i) \in (A \times G \times I)^n \mid \rho_S(T_{nt}(x, g, i) \circ \varepsilon_{s_n(a)}, T_t(x) \circ \sigma) \leq \epsilon\}) \\ \geq (\pi \times \bar{\gamma} \times \bar{\iota})^n(\tilde{A}_n(\epsilon) \cap (\tilde{B}_n(\delta) \times I'_n)),\end{aligned} \tag{A.60}$$

whereas the latter is, in view of (A.46), (A.54), and (A.55), greater than $1 - \epsilon$. ∎

Proof of Theorem 1: We use induction to show that, for each $\tau = 0, 1, ..., \bar{t} - t + 1$,

$$(\sigma_t \times \bar{\gamma}^\tau \times \bar{\iota}^\tau)^n(\tilde{A}_{n\tau}(\epsilon)) > 1 - \frac{\epsilon}{\bar{t} - t + 2}, \tag{A.61}$$

for any $\epsilon > 0$ and n large enough, where $\tilde{A}_{n\tau}(\epsilon) \in \mathcal{B}^n(S \times G^\tau \times I^\tau)$ is such that, for any $(s_t, g_{[t,t+\tau-1]}, i_{[t,t+\tau-1]}) \in \tilde{A}_{n\tau}(\epsilon)$,

$$\rho_S(T_{n,[t,t+\tau-1]}(x_{[t,t+\tau-1]}, g_{[t,t+\tau-1]}, i_{[t,t+\tau-1]}) \circ \varepsilon_{s_t}, T_{[t,t+\tau-1]}(x_{[t,t+\tau-1]}) \circ \sigma_t) < \epsilon. \tag{A.62}$$

Once the above is achieved, we can then define $\tilde{A}_n(\epsilon)$ required in the theorem by

$$\tilde{A}_n(\epsilon) = \bigcap_{\tau=0}^{\bar{t}-t+1} [\tilde{A}_{n\tau}(\epsilon) \times G^{n \cdot (\bar{t}-t+1-\tau)} \times I^{n \cdot (\bar{t}-t+1-\tau)}]. \tag{A.63}$$

This and (A.61) will lead to

$$(\sigma_t \times \bar{\gamma}^{\bar{t}-t+1} \times \bar{\iota}^{\bar{t}-t+1})^n(\tilde{A}_n(\epsilon)) > (1 - \frac{\epsilon}{\bar{t}-t+2})^{\bar{t}-t+2} > 1-\epsilon, \tag{A.64}$$

for any $\epsilon > 0$ and n large enough.

Now we proceed with the induction process. First, note that $T_{n,[t,t-1]} \circ \varepsilon_{s_t}$ is merely ε_{s_t} itself and $T_{[t,t-1]} \circ \sigma_t$ is merely σ_t itself. Hence, we will have (A.61) for $\tau = 0$ for any $\epsilon > 0$ and n large enough just by Lemma 1. Then, for some $\tau = 1, 2, ..., \bar{t} - t + 1$, suppose

$$(\sigma_t \times \bar{\gamma}^{\tau-1} \times \bar{\iota}^{\tau-1})^n(\tilde{A}_{n,\tau-1}(\epsilon)) > 1 - \frac{\epsilon}{\bar{t}-t+2}, \tag{A.65}$$

for any $\epsilon > 0$ and n large enough. We may apply Proposition 2 to the above, while at the same time identifying $S \times G^{\tau-1} \times I^{\tau-1}$ with A, $\sigma_t \times \bar{\gamma}^{\tau-1} \times \bar{\iota}^{\tau-1}$ with π, $x_{t+\tau-1}$ with x, $T_{n,[t,t+\tau-2]}(x_{[t,t+\tau-2]}, g_{[t,t+\tau-2]}, i_{[t,t+\tau-2]}) \circ \varepsilon_{s_t}$ with $\varepsilon_{s_n(a)}$, and $T_{[t,t+\tau-2]}(x_{[t,t+\tau-2]}) \circ \sigma_t$ with σ. This way, we will verify (A.61) for any $\epsilon > 0$ and n large enough. Therefore, the induction process can be completed. ∎

B Technical Developments in Section 6.2

Proof of Proposition 3: Because payoff functions are bounded, the value functions are bounded too. We then prove by induction on t. By (16), we know the result is true for $t = \bar{t}+1$. Suppose for some $t = \bar{t}, \bar{t}-1, ..., 2$, we have the continuity of $v_{t+1}(s_{t+1}, \sigma_{t+1}, x_{[t+1,\bar{t}]}, x_{t+1})$ in s_{t+1}. By this induction hypothesis, the probabilistic continuity of x_t, (S1), (F1), and the boundedness of the value functions, we see the continuity of the right-hand side of (17) in s_t. So, $v_t(s_t, \sigma_t, x_{[t\bar{t}]}, x_t)$ is continuous in s_t, and we have completed our induction process. ∎

Proof of Proposition 4: We prove by induction on t. By (16) and (23), we know the result is true for $t = \bar{t} + 1$. Suppose for some $t = \bar{t}, \bar{t} - 1, ..., 2$, we have the convergence of $v_{n,t+1}(s_{t+1,1}, \varepsilon_{s^n_{t+1,-1}}, x_{[t+1,\bar{t}]}, x_{t+1})$ to $v_{t+1}(s_{t+1,1}, \sigma_{t+1}, x_{[t+1,\bar{t}]}, x_{t+1})$ at an $s_{t+1,1}$-independent rate when $s_{t+1,-1} = (s_{t+1,2}, s_{t+1,3}, ...)$ is sampled from σ_{t+1}. Now, suppose $s_{t,-1} = (s_{t2}, s_{t3}, ...)$ is sampled from σ_t. Let also $g = (g_1, g_2, ...)$ be generated through sampling on $(G, \mathcal{B}(G), \bar{\gamma})$ and $i = (i_1, i_2, ...)$ be generated through sampling on $(I, \mathcal{B}(I), \bar{\iota})$. In the remainder of the proof, we let $s^n_t = (s_{t1}, s_{t2}, ..., s_{tn})$ for any arbitrary $s_{t1} \in S$, $g^n = (g_1, ..., g_n)$ and $i^n = (i_1, ..., i_n)$.

Due to Lemma 1, $\varepsilon_{s^n_{t,-1}}$ will converge to σ_t. By Lemma 2, $\varepsilon_{s^n_t}$ will converge to σ_t at an s_{t1}-independent rate. By Proposition 2, we know that $T_{nt}(x_t, g^n, i^n) \circ \varepsilon_{s^n_t}$ will converge to $T_t(x_t) \circ \sigma_t$ in probability at an s_{t1}-independent rate, and by Lemma 2 again, so will $[T_{nt}(x_t, g^n, i^n) \circ \varepsilon_{s^n_t}]_{-1}$ to $T_t(x_t) \circ \sigma_t$. Now Lemma 3 will lead to the convergence in probability of $\varepsilon_{s^n_{t,-1}, g^n_{-1}}$ to $\sigma_t \times \bar{\gamma}$. Due to x_t's probabilistic continuity, Lemma 4 will lead to the convergence in probability of $\varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}$ to $(\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}$. Thus,

1. $\tilde{f}_t(s_{t1}, x_t(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1})$ will converge to $\tilde{f}_t(s_{t1}, x_t(s_{t1}, g_1), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1})$ in probability at an s_{t1}-independent rate due to (F2);

2. $v_{n,t+1}(\tilde{s}_t(s_{t1}, x_t(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}, i_1), [T_{nt}(x_t, g^n, i^n) \circ \varepsilon_{s^n_t}]_{-1}, x_{[t+1,\bar{t}]}, x_{t+1})$ will converge to $v_{t+1}(\tilde{s}_t(s_{t1}, x_t(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}, i_1), T_t(x_t) \circ \sigma_t, x_{[t+1,\bar{t}]}, x_{t+1})$ in probability at an s_{t1}-independent rate due to the induction hypothesis; the latter will in turn converge to $v_{t+1}(\tilde{s}_t(s_{t1}, x_t(s_{t1}, g_1), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i_1), T(x_t) \circ \sigma_t, x_{[t+1,\bar{t}]}, x_{t+1})$ in probability at an s_{t1}-independent rate due to (S2) and Proposition 3.

As per-period payoffs are bounded, all value functions are bounded. The above convergences will then lead to the convergence of the right-hand side of (24) to the right-hand side of (17) at an s_{t1}-independent rate. That is, $v_{nt}(s_{t1}, \varepsilon_{s^n_{t,-1}}, x_{[t\bar{t}]}, x_t)$ will converge to $v_t(s_{t1}, \sigma_t, x_{[t\bar{t}]}, x_t)$ at a rate independent of s_{t1}. We have completed the induction process. ∎

Proof of Theorem 2: Let us consider subgames starting with some time $t = 1, 2, ..., \bar{t}$. For convenience, we let $\sigma_t = T_{[1,t-1]}(x^*_{[1,t-1]}) \circ \sigma_1$. Now let $s_t = (s_{t1}, s_{t2}, ...)$ be generated through sampling on $(S, \mathcal{B}(S), \sigma_t)$, $g = (g_1, g_2, ...)$ be generated through sampling on $(G, \mathcal{B}(G), \bar{\gamma})$, and $i = (i_1, i_2, ...)$ be generated through sampling on $(I, \mathcal{B}(I), \bar{\iota})$. In the remainder of the proof, we let $s^n_t = (s_{t1}, ..., s_{tn})$, $s^n_{t,-1} = (s_{t2}, ..., s_{tn})$, $g^n = (g_1, ..., g_n)$, and $i^n = (i_1, ..., i_n)$.

By Lemma 1 and Proposition 2, we know that $\varepsilon_{s^n_t} = \varepsilon_{(s_{t1}, ..., s_{tn})}$ converges to σ_t in probability, and also that $T_{nt}(x^*_t, g^n, i^n) \circ \varepsilon_{s^n_t}$ converges to $T_t(x^*_t) \circ \sigma_t$ in probability. Due to Lemma 2, $\varepsilon_{s^n_{t,-1}}$ and $[T_{nt}(x^*_t, g^n, i^n) \circ \varepsilon_{s^n_t}]_{-1}$ will have the same respective convergences. Also, Lemma 3 will lead to the convergence in probability of $\varepsilon_{s^n_{t,-1}, g^n_{-1}}$ to $\sigma_t \times \bar{\gamma}$. Due to x_t's probabilistic continuity, Lemma 4 will lead to the convergence in probability of $\varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}$ to $(\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}$. Then,

1. $\tilde{f}_t(s_{t1}, y(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g_{-1}} \cdot (1_S, x_t)^{-1})$ will converge to $\tilde{f}_t(s_{t1}, y(s_{t1}, g_1), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1})$ in probability at a y-independent rate due to (F2);

2. $v_{n,t+1}(\tilde{s}_t(s_{t1}, y(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}, i_1), [T_{nt}(x^*_t, g^n, i^n) \circ \varepsilon_{s^n_t}]_{-1}, x^*_{[t+1,\bar{t}]}, x^*_{t+1})$ will converge to $v_{t+1}(\tilde{s}_t(s_{t1}, y(s_{t1}, g_1), \varepsilon_{s^n_{t,-1}, g^n_{-1}} \cdot (1_S, x_t)^{-1}, i_1), T_t(x^*_t) \circ \sigma_t, x^*_{[t+1,\bar{t}]}, x^*_{t+1})$ in probability

at a y-independent rate due to Proposition 4, which will converge to $v_{t+1}(\tilde{s}_t(s_{t1}, y(s_{t1}, g_1), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i_1), T_t(x_t^*) \circ \sigma_t, x^*_{[t+1,\bar{t}]}, x^*_{t+1})$ in probability at a y-independent rate due to (S2) and Proposition 3.

As per-period payoffs are bounded, all value functions are bounded. By (17) and (24), the above convergences will then lead to the convergence of the left-hand side of (25) to the left-hand side of (19). At the same time, the right-hand side of (25) plus ϵ will converge to the right-hand side of (19) due to the convergence of $\varepsilon_{s^n_{t,-1}}$ to σ_t, Proposition 4, and the uniform boundedness of the value functions. By (19), as long as n is large enough, (25) will be true for any $\epsilon > 0$ and $y \in \mathcal{M}(S \times G, X)$. This would then lead to the final conclusion due to Theorem 1 and the boundedness of payoff functions. ∎

C Technical Developments in Section 7

Proof of Theorem 3: Let $\epsilon > 0$ be fixed. For $t = 1, 2, ...$ satisfying $t \geq \ln(6\bar{f}/(\epsilon \cdot (1 - \bar{\phi})))/\ln(1/\bar{\phi}) + 1$, we have from (37) and (38),

$$\mid v_{n\infty}(s_1, \varepsilon_{s_{-1}}, x^*, y) - v_{nt}(s_1, \varepsilon_{s_{-1}}, x^*, y) \mid < \frac{\epsilon}{6}. \tag{C.1}$$

Therefore, we need merely to select such a large t and show that, when n is large enough,

$$\int_{S^n} v_{nt}(s_1, \varepsilon_{s_{-1}}, x^*, x^*) \cdot (\sigma^*)^n(ds) \geq \int_{S^n} v_{nt}(s_1, \varepsilon_{s_{-1}}, x^*, y) \cdot (\sigma^*)^n(ds) - \frac{2\epsilon}{3}. \tag{C.2}$$

For $t = 1, 2, ...$, since (x^*, σ^*) forms an equilibrium for Γ, we know (34) is true. This, as well as (32) and (33), lead to

$$\bar{\phi}^{t-\tau} \cdot [\int_S v_\tau(s, \sigma^*, x^*, y) \cdot \sigma^*(ds) - \int_S v_\tau(s, \sigma^*, x^*, x^*) \cdot \sigma^*(ds)] \leq \frac{2\bar{\phi}^{t-1} \cdot \bar{f}}{1 - \bar{\phi}} \leq \frac{\epsilon}{3}. \tag{C.3}$$

for $\tau = 1, 2, ..., t$, $g \in G$, $s \in S$, and $y \in \mathcal{M}(S \times G, X)$.

We associate entities here with those defined in Section 4 when $\bar{t}$ there is fixed at the t here. To signify the difference in the two notational systems, we add superscript "K" to symbols defined for the previous section. For instance, we write v_τ^K for the v_τ defined in that section, which has a different meaning than the v_τ here. Now, our $\bar{\phi}^{t-\tau} \cdot v_\tau(s, \sigma^*, x^*, y)$ can be understood as $v^K_{t+1-\tau}(s, \sigma^*, x', y)$, with $x' = (x'_{t+1-\tau}, ..., x'_t) \in (\mathcal{M}(S \times G, X))^\tau$ being such that $x'_{t'} = x^*$ for $t' = t + 1 - \tau, ..., t$. Due to the association of σ^* with x^* through the

definition (30), we can understand σ^* as $T^K_{[1,\tau-1]}(x'_{[1,\tau-1]}) \circ \sigma^K_1$, where $x'_{[1,\tau-1]} = (x'_1, ..., x'_{\tau-1}) \in (\mathcal{M}(S \times G, X))^{\tau-1}$ is such that $x'_{t'} = x^*$ for $t' = 1, 2, ..., \tau - 1$.

With these correspondences, (C.3) can be translated into something akin to (19), with the only difference being that $-\epsilon/3$ should be added to all the right-hand sides. That is, we now know that the current (x^*, σ^*) offers an $(\epsilon/3)$-Markov equilibrium for the nonatomic game $\Gamma^{\mathcal{K}}(\sigma^*)$ with $\bar{t} = t$, $\tilde{s}^K_\tau = \tilde{s}$, and $\tilde{f}^K_\tau = \bar{\phi}^{\tau-1} \cdot \tilde{f}$. Even though Theorem 2 is nominally about going from an 0-equilibrium for the nonatomic game to ϵ-equilibria for finite games, we can follow exactly the same logic used to prove it to go from an $(\epsilon/3)$-equilibrium for the nonatomic game to $(2\epsilon/3)$-equilibria for finite games.

Thus, from one of the theorem's claims, we can conclude that, for n large enough and any $y \in \mathcal{M}(S \times G, X)$,

$$\int_{S^n} (\sigma^K_1)^n(ds) \cdot v^K_{nt}(s_1, \varepsilon_{s_{-1}}, x'_{[1t]}, x'_1) \geq \int_{S^n} (\sigma^K_1)^n(ds) \cdot v^K_{nt}(s_1, \varepsilon_{s_{-1}}, x'_{[1t]}, y) - \frac{2\epsilon}{3}, \tag{C.4}$$

where $x'_{[1t]}$ is again to be understood as the policy that takes action $x^*(s, g)$ whenever the most immediate state-shock pair is (s, g). But this translates into (C.2). ∎

D Technical Developments in Section 8.2

Use $\chi_{[1\bar{t}]} \equiv (\chi_t \mid t = 1, ..., \bar{t}) \in ((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ to denote a policy. For $t = 1, ..., \bar{t}$, suppose pre-action environment $\sigma_t \equiv (\sigma_{ts_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ has been determined. According to (9) and (10), the in-action environment $\mu_t \equiv (\mu_{ts_Ts_Ix})_{s_T=1,...,\bar{s}_T, s_I=1,...,s_I, x=1,...,\bar{x}} \in \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x})$, to be faced by all firms, will satisfy

$$\mu_{ts_Ts_Ix} = \sigma_{ts_Ts_I} \cdot \chi_{ts_Ts_Ix} \qquad \forall s_T = 1, ..., \bar{s}_T, s_I = 1, ..., \bar{s}_I, x = 1, ..., \bar{x}. \tag{D.1}$$

Combining this with (45), we see that the type-action marginal $\xi_t \equiv (\xi_{ts_Tx})_{s_T=1,...,\bar{s}_T, x=1,...,\bar{x}} = \mu_t|_{S_T \times X} \in \Delta(\bar{s}_T \cdot \bar{x})$ satisfies

$$\xi_{ts_Tx} = \sum_{s_I=1}^{\bar{s}_I} \mu_{ts_Ts_Ix} = \sum_{s_I=1}^{\bar{s}_I} \sigma_{ts_Ts_I} \cdot \chi_{ts_Ts_Ix} \qquad \forall s_T = 1, ..., \bar{s}_T, x = 1, ..., \bar{x}. \tag{D.2}$$

Let x_t be the element in $\mathcal{M}(S \times G, X)$ corresponding to the current $\chi_t \in (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ in the fashion of (48). From (12) and (47), we see that $\sigma_{t+1} \equiv (\sigma_{t+1,s_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} = T_t(x_t) \circ \sigma_t \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ should satisfy

$$\begin{aligned} \sigma_{t+1,s_Ts_I} = {} & \mathbf{1}(s_I = 0) \cdot \sigma_{ts_T0} + \mathbf{1}(s_I \geq 1) \cdot \sigma_{ts_Ts_I} \cdot \textstyle\sum_{x=1}^{\bar{x}} \chi_{ts_Ts_Ix} \cdot (1 - \tilde{\lambda}_{s_Tx}(\xi_t)) \\ & + \mathbf{1}(s_I \leq \bar{s}_I - 1) \cdot \sigma_{ts_T,s_I+1} \cdot \textstyle\sum_{x=1}^{\bar{x}} \chi_{ts_T,s_I+1,x} \cdot \tilde{\lambda}_{s_Tx}(\xi_t), \end{aligned} \tag{D.3}$$

for $s_T = 1, ..., \bar{s}_T$ and $s_I = 0, 1, ..., \bar{s}_I$. The above (D.3) shows that a firm's type will not change. It also reflects that, a firm will start period $t+1$ with s_I items when, either it starts period t with s_I items and whatever price choice x it makes does not bring in any demand or $s_I \leq \bar{s}_I - 1$, the firm starts period t with $s_I + 1$ items, and whatever price choice x it makes does bring in one unit of demand.

At a particular $t = 1, ..., \bar{t}$, we can understand (D.2) as $\xi_t \equiv (\xi_{ts_Tx})_{s_T=1,...,\bar{s}_T, x=1,...,\bar{x}} = Z^E_{FD,t}(\sigma_t, \chi_t)$ where $Z^E_{FD,t}$ is a mapping from $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1)) \times (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ to $\Delta(\bar{s}_T \cdot \bar{x})$, and understand (D.3) as $\sigma_{t+1} = Z^F_{FDE,t}(\sigma_t, \chi_t, \xi_t)$ where $Z^F_{FDE,t}$ is a mapping from $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1)) \times (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x})$ to $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$. Under the fixed game-defining $\sigma_1 \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$, let us now use $\xi_{[1\bar{t}]} \equiv (\xi_t | t = 1, ..., \bar{t}) = Z^E_D(\chi_{[1\bar{t}]} \mid \sigma_1)$, where $Z^E_D(\cdot|\sigma_1)$ is a mapping from $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ to $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$, to denote the attainment of the ξ_t's through the iterations $\xi_t = Z^E_{FD,t}(\sigma_t, \chi_t)$ and $\sigma_{t+1} = Z^F_{FDE,t}(\sigma_t, \chi_t, \xi_t)$ over $t = 1, ..., \bar{t}$.

Suppose a type-price-distribution sequence $\xi_{[1\bar{t}]} \equiv (\xi_t \mid t = 1, ..., \bar{t}) \in (\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ is given. Let $v^*_{ts_Ts_I}$ be the optimal total expected payoff a type-s_T firm can make from period t to $\bar{t}$ when it starts with inventory level $s_I = 0, 1, ..., \bar{s}_I$. As a terminal condition, we have

$$v^*_{\bar{t}+1,s_Ts_I} = 0, \qquad \forall s_T = 0, 1, ..., \bar{s}_T, s_I = 1, ..., \bar{s}_I. \tag{D.4}$$

For $t = \bar{t}, \bar{t} - 1, ..., 1$ and $s_T = 1, ..., \bar{s}_T$, we have, due to (46) and (47), $v^*_{ts_T0} = 0$; also, for $s_I = 1, ..., \bar{s}_I$, we have $v^*_{ts_Ts_I}$ equal to

$$-\bar{h}_{s_T} \cdot s_I + \sup_{\chi \in \partial\Delta(\bar{x})} \left\{ \sum_{x=1}^{\bar{x}} \chi_x \cdot [\tilde{\lambda}_{s_Tx}(\xi_t) \cdot (\bar{p}_x + v^*_{t+1,s_T,s_I-1}) + (1 - \tilde{\lambda}_{s_Tx}(\xi_t)) \cdot v^*_{t+1,s_Ts_I}] \right\}. \tag{D.5}$$

The above (D.5) allows any type-s_T firm to choose the best probabilities χ_x to use on price choices x; when a particular x has been settled on, there will be a $\tilde{\lambda}_{s_Tx}(\xi_t)$ chance for the firm to earn $\bar{p}_x$ in revenue and enter period $t+1$ with one less item in stock, and also a $1 - \tilde{\lambda}_{s_Tx}(\xi_t)$ chance for it to earn nothing and enter period $t+1$ with the same stock level as it had to start period t with. Here, the impact of the current individual firm's decision on the overall environment's future evolution is ignored—ξ_t is given; yet, the future evolution of the current firm's own inventory level is a concern in its decision making.

For $t = \bar{t} + 1, \bar{t}, ..., 1$, define $v^*_t \equiv (v^*_{ts_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I} \in \Re^{\bar{s}_T \cdot \bar{s}_I}$. At a particular $t = \bar{t}, \bar{t}-1, ..., 1$, with the understanding that $v^*_{t+1,s_T0} = 0$ for any $s_T = 1, ..., \bar{s}_T$, we can write (D.5) for any particular $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$ as $v^*_{ts_Ts_I} = Z^V_{EV,ts_Ts_I}(\xi_t, v^*_{t+1})$, where Z^V_{EV,ts_Ts_I} is a mapping from $\Delta(\bar{s}_T \cdot \bar{x}) \times \Re^{\bar{s}_T \cdot \bar{s}_I}$ to $\Re$. At the same time, we can use $\chi_{ts_Ts_I} \in \mathbf{Z}^D_{EV,ts_Ts_I}(\xi_t, v^*_{t+1})$ to denote $\chi_{ts_Ts_I}$'s optimality to (D.5), where $\mathbf{Z}^D_{EV,ts_Ts_I}$ is a correspondence

from $\Delta(\bar{s}_T \cdot \bar{x}) \times \Re^{\bar{s}_T \cdot \bar{s}_I}$ to $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$. Combining these over all the $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$, we can write $v_t^* = Z^V_{EV,t}(\xi_t, v_{t+1}^*)$ where $Z^V_{EV,t} = \prod_{s_T=1}^{\bar{s}_T} \prod_{s_I=1}^{\bar{s}_I} Z^V_{EV,ts_Ts_I}$ and $\chi_t \equiv (\chi_{ts_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I} \in \mathbf{Z}^D_{EV,t}(\xi_t, v_{t+1}^*)$ where $\mathbf{Z}^D_{EV,t} = \prod_{s_T=1}^{\bar{s}_T} \prod_{s_I=1}^{\bar{s}_I} \mathbf{Z}^D_{EV,ts_Ts_I}$. Suppose we start $v^*_{\bar{t}+1}$ with the all-zero vector which satisfies (D.4), and for $t = \bar{t}, \bar{t}-1, ..., 1$ obtain v_t^* by applying $v_t^* = Z^V_{EV,t}(\xi_t, v_{t+1}^*)$. Then, we can define a correspondence $\mathbf{Z}^D_E$ from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ so that $\chi_{[1\bar{t}]} \equiv (\chi_t | t = 1, ..., \bar{t}) \in \mathbf{Z}^D_E(\xi_{[1\bar{t}]})$ if and only if $\chi_t \in \mathbf{Z}^D_{EV,t}(\xi_t, v^*_{t+1})$ for every $t = \bar{t}, \bar{t}-1, ..., 1$.

Equate correspondence $\mathbf{Z}^D_D(\cdot|\sigma_1)$ on $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ with $\mathbf{Z}^D_E \circ Z^E_D(\cdot|\sigma_1)$, so that

$$\chi'_{[1\bar{t}]} \in \mathbf{Z}^D_D(\chi_{[1\bar{t}]}|\sigma_1) \qquad \text{if and only if} \qquad \chi'_{[1\bar{t}]} \in \mathbf{Z}^D_E(Z^E_D(\chi_{[1\bar{t}]}|\sigma_1)). \tag{D.6}$$

Due to definitions and arguments made from (16) to (20) and the correspondence established at (48), we know that a fixed point $\chi^*_{[1\bar{t}]}$ satisfying $\chi^*_{[1\bar{t}]} \in \mathbf{Z}^D_D(\chi^*_{[1\bar{t}]}|\sigma_1)$ will provide a Markov equilibrium $x^*_{[1\bar{t}]} \in (\mathcal{M}(S \times G, X))^{\bar{t}}$ to the nonatomic game $\Gamma(\sigma_1)$.

Proof of Proposition 5: Let us use Kakutani's Theorem to prove existence of the fixed points. Since $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ is a non-empty, compact, and convex subset of the Euclidean space $\Re^{\bar{s}_T \cdot \bar{s}_I \cdot \bar{x} \cdot \bar{t}}$, it will suffice if $\mathbf{Z}^D_D(\chi_{[1\bar{t}]}|\sigma_1)$ is non-empty and convex at every $\chi_{[1\bar{t}]} \in ((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ and as a correspondence, $\mathbf{Z}^D_D(\cdot|\sigma_1)$ is closed.

From (D.2), we see that each $Z^E_{FD,t}$ is a continuous mapping from $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1)) \times (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ to $\Delta(\bar{s}_T \cdot \bar{x})$. From (P1) and (D.3), we see that each $Z^F_{FDE,t}$ is a continuous mapping from $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1)) \times (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x})$ to $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$. When used iteratively, these will lead to the continuity of $Z^E_D(\cdot|\sigma_1)$ as a mapping from $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ to $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$. Indeed, the dependence on $\sigma_1 \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ is also continuous, although this property is not used here.

For any $t = \bar{t}+1, \bar{t}, ..., 1$, $s_T = 1, ..., \bar{s}_T$, and $s_I = 0, 1, ..., \bar{s}_I$, define Z^V_{E,ts_Ts_I} from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $\Re$ as follows. First, let $Z^V_{E,ts_Ts_I}(\xi_{[1\bar{t}]}) = 0$ whenever $t = \bar{t}+1$ or $s_I = 0$. Then, for $t = \bar{t}, \bar{t}-1, ..., 1$, let

$$Z^V_{E,ts_Ts_I}(\xi_{[1\bar{t}]}) = -\bar{h}_{s_T} \cdot s_I + \sup_{\chi \in \partial\Delta(\bar{x})} \left\{ \sum_{x=1}^{\bar{x}} \chi_x \cdot W_{ts_Ts_Ix}(\xi_{[1\bar{t}]}) \right\}, \tag{D.7}$$

for any $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$, where each

$$W_{ts_Ts_Ix}(\xi_{[1\bar{t}]}) = \tilde{\lambda}_{s_Tx}(\xi_t) \cdot (\bar{p}_x + Z^V_{E,t+1,s_T,s_I-1}(\xi_{[1\bar{t}]})) + (1 - \tilde{\lambda}_{s_Tx}(\xi_t)) \cdot Z^V_{E,t+1,s_Ts_I}(\xi_{[1\bar{t}]}). \tag{D.8}$$

By their definitions, we know that each $Z^V_{E,\bar{t}+1,s_Ts_I}$ is a continuous mapping from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $\Re$; so is each Z^V_{E,ts_T0}. Suppose for some $t = \bar{t}, \bar{t}-1, ..., 1$, each $Z^V_{E,t+1,s_Ts_I}$ for $s_T = 1, ..., \bar{s}_T$

and $s_I = 1, ..., \bar{s}_I$ is a continuous mapping from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $\Re$. Then by (P1) and (D.8), we know that $W_{ts_Ts_Ix}$ for every $s_T = 1, ..., \bar{s}_T$, $s_I = 1, ..., \bar{s}_I$, and $x = 1, ..., \bar{x}$ is a continuous mapping from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $\Re$. Now in the optimization problem of (D.7), the objective is continuous in χ and $\xi_{[1\bar{t}]}$, while the feasible region is compact as a set of the χ's and even independent of $\xi_{[1\bar{t}]}$. So by Berge's Theorem, we can obtain the continuity of Z^V_{E,ts_Ts_I} as a mapping from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $\Re$ at every $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$. An induction process has been completed.

Comparing (D.5) with (D.7) and (D.8), we see that the $v^*_{ts_Ts_I}$'s defined for a given $\xi_{[1\bar{t}]}$ are exactly the same as the $Z^V_{E,ts_Ts_I}(\xi_{[1\bar{t}]})$'s. This way, $\chi_{[1\bar{t}]} \equiv ((\chi_{ts_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I} | t = 1, ..., \bar{t}) \in \mathbf{Z}^D_E(\xi_{[1\bar{t}]})$ if and only if $\chi_{ts_Ts_I} \in \mathbf{Z}^D_{E,ts_Ts_I}(\xi_{[1\bar{t}]})$ for every $t = \bar{t}, \bar{t}-1, ..., 1$, $s_T = 1, ..., \bar{s}_T$, and $s_I = 1, ..., \bar{s}_I$, where

$$\mathbf{Z}^D_{E,ts_Ts_I}(\xi_{[1\bar{t}]}) = \text{argmax}_{\chi \in \partial\Delta(\bar{x})} \left\{ \sum_{x=1}^{\bar{x}} \chi_x \cdot W_{ts_Ts_Ix}(\xi_{[1\bar{t}]}) \right\}. \tag{D.9}$$

Since $W_{ts_Ts_Ix}$ and hence the objective of the optimization problem in (D.9) is continuous in $\xi_{[1\bar{t}]}$, the latter being also jointly continuous in χ, and the feasible region is compact as a set of the χ's and also independent of $\xi_{[1\bar{t}]}$, we know from Berge's Theorem that $\mathbf{Z}^D_{E,ts_Ts_I}(\xi_{[1\bar{t}]})$ is non-empty at every $\xi_{[1\bar{t}]}$ and that $\mathbf{Z}^D_{E,ts_Ts_I}$ as a correspondence is closed. Since the objective in the optimization problem of (D.9) is affine in $\chi \equiv (\chi_x)_{x=1,...,\bar{x}}$ and the feasible region is convex, we also know that every $\mathbf{Z}^D_{E,ts_Ts_I}(\xi_{[1\bar{t}]})$ is convex.

In combination, we have that $\mathbf{Z}^D_E(\xi_{[1\bar{t}]})$ is non-empty and convex at each $\xi_{[1\bar{t}]}$, and that $\mathbf{Z}^D_E$ as a correspondence from $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$ to $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ is closed. Combine this further with the continuity of $Z^E_D(\cdot|\sigma_1)$ as a mapping from $((\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I})^{\bar{t}}$ to $(\Delta(\bar{s}_T \cdot \bar{x}))^{\bar{t}}$, and we can assert that $\mathbf{Z}^D_D(\cdot|\sigma_1) = \mathbf{Z}^D_E \circ Z^E_D(\cdot|\sigma_1)$ as defined in (D.6) has the desired properties. ∎

E Technical Developments in Section 8.3

Let $\chi \in (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ denote an action plan. When given pre-action environment $\sigma \equiv (\sigma_{s_Ts_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ for the period, the in-action environment $\mu \in \Delta(\bar{s}_T \cdot \bar{s}_I \cdot \bar{x})$ is still given by (D.1), whose marginal $\xi = \mu|_{S_T \times X}$ is again defined by (D.2). Let x be the element in $\mathcal{M}(S \times G, X)$ that corresponds to the current χ through (48). Due

to (29) and (50), $\sigma' \equiv (\sigma'_{s_T s_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} = T(x) \circ \sigma \in \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ should satisfy

$$\begin{aligned}\sigma'_{s_T s_I} = \mathbf{1}(s_I = \bar{s}_I) \cdot \sigma_{s_T 0} + \mathbf{1}(s_I \geq 1) \cdot \sigma_{s_T s_I} \cdot \textstyle\sum_{x=1}^{\bar{x}} \chi_{s_T s_I x} \cdot (1 - \tilde{\lambda}_{s_T x}(\xi)) \\ + \mathbf{1}(s_I \leq \bar{s}_I - 1) \cdot \sigma_{s_T, s_I+1} \cdot \textstyle\sum_{x=1}^{\bar{x}} \chi_{s_T, s_I+1, x} \cdot \tilde{\lambda}_{s_T x}(\xi),\end{aligned} \tag{E.1}$$

for $s_T = 1, ..., \bar{s}_T$ and $s_I = 0, 1, ..., \bar{s}_I$. Besides the same last two terms in (D.3), the above (E.1) has a different first term reflecting the scheduled production that always restores the inventory level back to $\bar{s}_I$ from 0. Due to its stationarity, we can understand (D.2) as $\xi = Z^E_{FD}(\sigma, \chi)$. Also, denote (E.1) by $\sigma' = Z^F_{FDE}(\sigma, \chi, \xi)$ where Z^F_{FDE} is a mapping from $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1)) \times (\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x})$ to $\partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$.

Given joint type-price distribution $\xi \in \Delta(\bar{s}_T \cdot \bar{x})$ and vector $v \equiv (v_{s_T s_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} \equiv \Re^{\bar{s}_T \cdot (\bar{s}_I+1)}$, we are prompted by (49) and (50) to define

$$v'_{s_T 0} = -\bar{k}_{s_T} + \bar{\phi} \cdot v_{s_T \bar{s}_I}, \qquad \forall s_T = 1, ..., \bar{s}_T, \tag{E.2}$$

and for $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$, define $v'_{s_T s_I}$ as

$$-\bar{h}_{s_T} \cdot s_I + \sup_{\chi \in \partial\Delta(\bar{x})} \left\{ \sum_{x=1}^{\bar{x}} \chi_x \cdot [\tilde{\lambda}_{s_T x}(\xi) \cdot (\bar{p}_x + \bar{\phi} \cdot v_{s_T, s_I - 1}) + (1 - \tilde{\lambda}_{s_T x}(\xi)) \cdot \bar{\phi} \cdot v_{s_T s_I}] \right\}. \tag{E.3}$$

(E.2) reflects that, at the expense of $\bar{k}_{s_T}$, a type-s_T firm's inventory level will be automatically restored to $\bar{s}_I$ once it has reached 0. Meanwhile, (E.3) is almost the same as (D.5), except that the current version takes into account the discount factor $\bar{\phi}$. We can view (E.2) and (E.3) as $v' = Z^V_{EV}(\xi, v)$, where Z^V_{EV} is a mapping from $\Delta(\bar{s}_T \cdot \bar{x}) \times \Re^{\bar{s}_T \cdot (\bar{s}_I+1)}$ to $\Re^{\bar{s}_T \cdot (\bar{s}_I+1)}$.

Lemma 5 *At each fixed $\xi \in \Delta(\bar{s}_T \cdot \bar{x})$, the operator $Z^V_{EV}(\xi, \cdot)$ has a unique fixed point $Z^V_E(\xi)$. Also, the fixed-point operator Z^V_E is continuous on $\Delta(\bar{s}_T \cdot \bar{x})$.*

Proof: For $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$, define mapping $Z^V_{DEV, s_T s_I}$ from $\partial\Delta(\bar{x}) \times \Delta(\bar{s}_T \cdot \bar{x}) \times \Re^{\bar{s}_T \cdot (\bar{s}_I+1)}$ to $\Re$, so that

$$Z^V_{DEV, s_T s_I}(\chi, \xi, v) = \sum_{x=1}^{\bar{x}} \chi_x \cdot [\tilde{\lambda}_{s_T x}(\xi) \cdot (\bar{p}_x + \bar{\phi} \cdot v_{s_T, s_I - 1}) + (1 - \tilde{\lambda}_{s_T x}(\xi)) \cdot \bar{\phi} \cdot v_{s_T s_I}]. \tag{E.4}$$

For any $u \equiv (u_{s_T s_I})_{s_T=1,...,\bar{s}_T, s_I=0,1,...,\bar{s}_I} \in \Re^{\bar{s}_T \cdot (\bar{s}_I+1)}$, let $||u|| = \max_{s_T=1}^{\bar{s}_T} \max_{s_I=0}^{\bar{s}_I} |u_{s_T s_I}|$. We can check that

$$|Z^V_{DEV, s_T s_I}(\chi, \xi, v) - Z^V_{DEV, s_T s_I}(\chi, \xi, v')| \leq \bar{\phi} \cdot ||v - v'||. \tag{E.5}$$

For $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$, define mapping $Z^V_{EV,s_T s_I}$ from $\Delta(\bar{s}_T \cdot \bar{x}) \times \Re^{\bar{s}_T \cdot (\bar{s}_I + 1)}$ to $\Re$, so that

$$Z^V_{EV,s_T s_I}(\xi, v) = \sup_{\chi \in \partial\Delta(\bar{x})} Z^V_{DEV,s_T s_I}(\chi, \xi, v). \tag{E.6}$$

Suppose, without loss of generality, that $Z^V_{EV,s_T s_I}(\xi, v) < Z^V_{EV,s_T s_I}(\xi, v')$. Let $\chi' \in \partial\Delta(\bar{x})$ be such that $Z^V_{EV,s_T s_I}(\xi, v') = Z^V_{DEV,s_T s_I}(\chi', \xi, v')$. Then by (E.5) and (E.6),

$$Z^V_{EV,s_T s_I}(\xi, v) \geq Z^V_{DEV,s_T s_I}(\chi', \xi, v) \geq Z^V_{DEV,s_T s_I}(\chi', \xi, v') - \bar{\phi} \cdot ||v - v'||, \tag{E.7}$$

which is equal to $Z^V_{EV,s_T s_I}(\xi, v') - \bar{\phi} \cdot ||v - v'||$. So by symmetry,

$$|Z^V_{EV,s_T s_I}(\xi, v) - Z^V_{EV,s_T s_I}(\xi, v')| \leq \bar{\phi} \cdot ||v - v'||. \tag{E.8}$$

Combining (E.2) to (E.8), we see that

$$||Z^V_{EV}(\xi, v) - Z^V_{EV}(\xi, v')|| \leq \bar{\phi} \cdot ||v - v'||. \tag{E.9}$$

In other words, Z^V_{EV} is a contraction mapping on the Banach space $\Re^{\bar{s}_T \cdot (\bar{s}_I + 1)}$ equipped with the $||\cdot||$-norm. Thus, it has a unique fixed point $Z^V_E(\xi)$ satisfying $Z^V_E(\xi) = Z^V_{EV}(\xi, Z^V_E(\xi))$. This point also happens to be $\lim_{n \to +\infty}(Z^V_{EV}(\xi, \cdot))^n(v_0)$ for any $v_0 \in \Re^{\bar{s}_T \cdot (\bar{s}_I + 1)}$.

At the same time, due to (P1) and (E.4), each $Z^V_{DEV,s_T s_I}$ is continuous in ξ. Let $||\cdot||$ also stand for the supremum norm of $\Re^{\bar{s}_T \cdot \bar{x}}$. Using arguments similar to those employed from (E.5) to (E.9), we can establish that, for any fixed $\xi \in \Delta(\bar{s}_T \cdot \bar{x})$ and any $\epsilon > 0$, there is $\delta > 0$, so that for any $\xi' \in \Delta(\bar{s}_T \cdot \bar{x})$ satisfying $||\xi - \xi'|| < \delta$,

$$||Z^V_{EV}(\xi, Z^V_E(\xi)) - Z^V_{EV}(\xi', Z^V_E(\xi))|| < \epsilon. \tag{E.10}$$

Consequently,

$$\begin{aligned}
&||Z^V_E(\xi) - Z^V_E(\xi')|| = ||Z^V_E(\xi) - \lim_{n \to +\infty}(Z^V_{EV}(\xi', \cdot))^n(Z^V_E(\xi))|| \\
&\quad \leq \textstyle\sum_{n=1}^{+\infty} ||(Z^V_{;EV}(\xi', \cdot))^{n-1}(Z^V_E(\xi)) - (Z^V_{EV}(\xi', \cdot))^n(Z^V_E(\xi))|| \\
&\quad = \textstyle\sum_{n=1}^{+\infty} ||(Z^V_{EV}(\xi', \cdot))^{n-1}(Z^V_{EV}(\xi, Z^V_E(\xi))) - (Z^V_{EV}(\xi', \cdot))^{n-1}(Z^V_{EV}(\xi', Z^V_E(\xi)))|| \\
&\quad \leq \textstyle\sum_{n=1}^{+\infty} \bar{\phi}^{n-1} \cdot \epsilon = \epsilon/(1 - \bar{\phi}),
\end{aligned} \tag{E.11}$$

where the first equality utilizes the fact that $Z^V_E(\xi') = \lim_{n \to +\infty}(Z^V_{EV}(\xi', \cdot))^n(Z^V_E(\xi))$, the first inequality is obtainable from a routine expansion, the second equality exploits the fact that $Z^V_E(\xi) = Z^V_{EV}(\xi, Z^V_E(\xi))$, the second inequality follows from combined uses of (E.9) and (E.10), and the last equality is an identity. ∎

Following the mapping Z^V_E, let us define correspondence $\mathbf{Z}^D_E$ from $\Delta(\bar{s}_T \cdot \bar{x})$ to $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$, so that $\chi \equiv (\chi_{s_T s_I})_{s_T=1,...,\bar{s}_T, s_I=1,...,\bar{s}_I} \in \mathbf{Z}^D_E(\xi)$ if and only if each $\chi_{s_T s_I}$ maximizes

$$\sum_{x=1}^{\bar{x}} \chi_x \cdot [\tilde{\lambda}_{s_T x}(\xi) \cdot (\bar{p}_x + \bar{\phi} \cdot Z^V_{E,s_T,s_I-1}(\xi)) + (1 - \tilde{\lambda}_{s_T x}(\xi)) \cdot \bar{\phi} \cdot Z^V_{E,s_T s_I}(\xi)]. \tag{E.12}$$

Define correspondence $\mathbf{Z}^{DEF}_{DEF}$ on $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x}) \times \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$, so that

$$\begin{aligned} (\chi', \xi', \sigma') \in \mathbf{Z}^{DEF}_{DEF}(\chi, \xi, \sigma) \qquad & \text{if and only if} \\ \chi' \in \mathbf{Z}^D_E(\xi'), \qquad \xi' = Z^E_{FD}(\sigma', \chi), \qquad & \sigma' = Z^F_{FDE}(\sigma, \chi, \xi). \end{aligned} \tag{E.13}$$

Suppose $(\chi^*, \xi^*, \sigma^*)$ is a fixed point for $\mathbf{Z}^{DEF}_{DEF}$, so that $(\chi^*, \xi^*, \sigma^*) \in \mathbf{Z}^{DEF}_{DEF}(\chi^*, \xi^*, \sigma^*)$. The fact that $\sigma^* = Z^F_{FDE}(\sigma^*, \chi^*, \xi^*) = Z^F_{FDE}(\sigma^*, \chi^*, Z^E_{FD}(\sigma^*, \chi^*))$ indicates that σ^* is associated with x^* in the sense of (30), where x^* is the element of $\mathcal{M}(S \times G, X)$ that corresponds to χ^* through (48). Meanwhile, $\chi^* \in \mathbf{Z}^D_E(\xi^*)$ implies the validity of (34). Therefore, $(\chi^*, \xi^*, \sigma^*)$ will provide a Markov equilibrium $x^* \in \mathcal{M}(S \times G, X)$ to the stationary nonatomic game Γ.

Proof of Proposition 6: We still use Kakutani's Theorem to prove existence of the fixed points. Since $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I} \times \Delta(\bar{s}_T \cdot \bar{x}) \times \partial\Delta(\bar{s}_T \cdot (\bar{s}_I + 1))$ is a non-empty, compact, and convex subset of the Euclidean space $\Re^{\bar{s}_T + \bar{s}_T \cdot \bar{s}_I + \bar{s}_T \cdot \bar{x} + \bar{s}_T \cdot \bar{s}_I \cdot \bar{x}}$, it will suffice if $\mathbf{Z}^{DEF}_{DEF}(\chi, \xi, \sigma)$ is non-empty and convex at every (χ, ξ, σ) and as a correspondence, $\mathbf{Z}^{DEF}_{DEF}$ is closed.

From the proof of Proposition 5, we know that Z^E_{FD} and Z^F_{FDE} are continuous mappings. For $s_T = 1, ..., \bar{s}_T$, $s_I = 1, ..., \bar{s}_I$, and $x = 1, ..., \bar{x}$, let

$$W_{s_T s_I x}(\xi) = \tilde{\lambda}_{s_T x}(\xi) \cdot (\bar{p}_x + Z^V_{E,s_T,s_I-1}(\xi)) + (1 - \tilde{\lambda}_{s_T x}(\xi)) \cdot Z^V_{E,s_T s_I}(\xi). \tag{E.14}$$

The definition (E.12) dictates that, for any $s_T = 1, ..., \bar{s}_T$ and $s_I = 1, ..., \bar{s}_I$,

$$\mathbf{Z}^D_{E,s_T s_I}(\xi) = \text{argmax}_{\chi \equiv (\chi_x)_{x=1,..,\bar{x}} \in \partial\Delta(\bar{x})} \left\{ \sum_{x=1}^{\bar{x}} \chi_x \cdot W_{s_T s_I x}(\xi) \right\}. \tag{E.15}$$

By Lemma 5 and (E.14), each $W_{s_T s_I x}$ and hence the objective of the optimization problem in (E.15) is continuous in ξ, the latter being also jointly continuous in χ. Meanwhile, the feasible region of the same problem is compact as a set of the χ's and also independent of ξ. So by Berge's Theorem, $\mathbf{Z}^D_{E,s_T s_I}(\xi)$ is non-empty at every ξ and that $\mathbf{Z}^D_{E,s_T s_I}$ as a correspondence is closed. Since the objective in the optimization problem of (E.15) is affine in $\chi \equiv (\chi_x)_{x=1,...,\bar{x}}$ and the feasible region is convex, every $\mathbf{Z}^D_{E,s_T s_I}(\xi)$ is also convex.

In combination, we have that $\mathbf{Z}^D_E(\xi)$ is non-empty and convex at each ξ, and that $\mathbf{Z}^D_E$ as a correspondence from $\Delta(\bar{s}_T \cdot \bar{x})$ to $(\partial\Delta(\bar{x}))^{\bar{s}_T \cdot \bar{s}_I}$ is closed. Combine this further with

the continuities of Z^E_{FD} and Z^F_{FDE}, and we can assert that $\mathbf{Z}^{DEF}_{DEF}$ as defined in (E.13) from $(\partial\Delta(\bar{x}))^{\bar{s}_T\cdot\bar{s}_I}\times\Delta(\bar{s}_T\cdot\bar{x})\times\partial\Delta(\bar{s}_T\cdot(\bar{s}_I+1))$ to itself has the desired properties. ∎

F Technical Developments in Section 9.2

Proof of Proposition 7: First, suppose action plan $\chi_t(s)$ is generated from a given shock-driven $x_t(s,g)$ in the fashion of (6). Then, for any $S'\in\mathcal{B}(S)$ and $X'\in\mathcal{B}(X)$, we can follow (54) to obtain

$$\begin{array}{ll}(\sigma\otimes\chi_t)(S'\times X') & =\int_{S'}\bar{\gamma}((x_t(s,\cdot))^{-1}(X'))\cdot\sigma_t(ds)\\ & =\int_S\int_G\mathbf{1}[(s,x_t(s,g))\in S'\times X']\cdot\bar{\gamma}(dg)\cdot\sigma_t(ds).\end{array}\tag{F.1}$$

For discrete spaces S and X, this is the same as $\sigma_t\otimes\chi_t$ being equal to the $(\sigma_t\times\bar{\gamma})\cdot(1_S,x_t)^{-1}$ that is defined through (9) and (10).

Second, when $x_t(s,g)$ is generated from $\chi_t(s)$ in the fashion of (8), we can derive that $(\sigma_t\times\bar{\gamma})\cdot(1_S,x_t)^{-1}$ is equal to $\sigma_t\otimes\chi_t$. Indeed, for any $S'\in\mathcal{B}(S)$ and $x\in\Re$,

$$\begin{array}{l}\int_S\int_G\mathbf{1}[(s,x_t(s,g))\in S'\times((-\infty,x]\cap X)]\cdot\bar{\gamma}(dg)\cdot\sigma_t(ds)\\ =\int_{S'}\bar{\gamma}(\{g\in G|x_t(s,g)\in(-\infty,x]\cap X\})\cdot\sigma_t(ds)=\int_{S'}\sigma_t(ds)\cdot\chi_t(s|(-\infty,x]\cap X),\end{array}\tag{F.2}$$

where the second equality utilizes both (8) and the fact that $\bar{\gamma}$ is uniform on $[0,1]$. ∎

Proof of Proposition 8: First, suppose action plan $\chi_t(s)$ is generated from a given shock-driven $x_t(s,g)$ in the fashion of (6) and the state transition rule $\tilde{g}_t(s,x,\mu)$ is generated from a given shock-driven $\tilde{s}_t(s,x,\mu,i)$ in the fashion of (3). Then, for any $\sigma_t\in\mathcal{P}(S)$ and $S'\in\mathcal{B}(S)$, we can follow Proposition 7 and (55) to obtain

$$\begin{array}{l}[\hat{T}_t(\chi_t)\circ\sigma_t](S')=\int_S\sigma_t(ds)\cdot\int_X[\bar{\gamma}\cdot(x_t(s,\cdot))^{-1}](dx)\times\\ \qquad\qquad\times\bar{\iota}(\{i\in I|\tilde{s}_t(s,x,(\sigma_t\times\bar{\gamma})\cdot(1_S,x_t)^{-1},i)\in S'\})\\ =\int_S\int_G\int_I\mathbf{1}[\tilde{s}_t(s,x_t(s,g),(\sigma_t\times\bar{\gamma})\cdot(1_S,x_t)^{-1},i)\in S']\cdot\bar{\iota}(di)\cdot\bar{\gamma}(dg)\cdot\sigma_t(ds).\end{array}\tag{F.3}$$

This amounts to $\hat{T}_t(\chi_t)\circ\sigma_t$ being equal to $T_t(x_t)\circ\sigma_t$ as defined through (12) and (13).

Second, when $x_t(s,g)$ is generated from $\chi_t(s)$ in the fashion of (8) and $\tilde{s}_t(s,x,\mu,i)$ is generated from $\tilde{g}_t(s,x,\mu)$ in the fashion of (5), we can derive that $T_t(x_t)$ is equal to $\hat{T}_t(\chi_t)$.

Indeed, for any $\sigma_t \in \mathcal{P}(S)$ and $s \in \Re$,

$$\begin{aligned}
&\int_S \int_G \int_I \mathbf{1}[\tilde{s}_t(s, x_t(s,g), (\sigma_t \times \bar{\gamma}) \cdot (1_S, x_t)^{-1}, i) \in (-\infty, s) \cap S] \cdot \bar{\iota}(di) \cdot \bar{\gamma}(dg) \cdot \sigma_t(ds) \\
&\quad = \int_S \sigma_t(ds) \cdot \int_X \chi_t(s|dx) \cdot \bar{\iota}(\{i \in I | \tilde{s}_t(s, x, \sigma_t \otimes \chi_t, i) \in (-\infty, s] \cap S\}) \\
&\quad = \int_S \sigma_t(ds) \cdot \int_X \chi_t(s|dx) \cdot \tilde{g}_t(s, x, \sigma_t \otimes \chi_t | (-\infty, s] \cap S),
\end{aligned} \tag{F.4}$$

where the first equality is due to Proposition 7 and (8), and the second equality utilizes both (5) and the fact that $\bar{\iota}$ is uniform on $[0, 1]$. ∎

References

[1] Adlakha, S. and R. Johari. 2013. Mean Field Equilibrium in Dynamic Games with Complementarities. *Operations Research*, **61**, pp. 971-989.

[2] Adlakha, S., R. Johari, and G.Y. Weintraub. 2011. Equilibria of Dynamic Games with Many Players: Existence, Approximation, and Market Structure. Working Paper, California Institute of Technology.

[3] Al-Najjar, N.I. and R. Smorodinsky. 2001. Large Nonanonymous Repeated Games. *Games and Economic Behavior*, **37**, pp. 26-39.

[4] Al-Najjar, N.I. 2008. Large Games and the Law of Large Numbers. *Games and Economic Behavior*, **64**, pp. 1-34.

[5] Aumann, R.J. 1964. Markets with a Continuum of Traders. *Econometrica*, **32**, pp. 39-50.

[6] Bergin, J. and D. Bernhardt. 1995. Anonymous Sequential Games: Existence and Characterization of Equilibria. *Economic Theory*, **5**, pp. 461-489.

[7] Carmona, G. 2004. Nash Equilibria of Games with a Continuum of Players. Working Paper, Universidade Nova de Lisboa.

[8] Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan. 1994. Stationary Markov Equilibria. *Econometrica*, **62**, pp. 745-781.

[9] Ethier, S.N. and T.G. Kurtz. 1986. *Markov Processes: Characterization and Convergence*. John Wiley & Sons, New York.

[10] Feldman, M. and C. Gilles. 1985. An Expository Note on Individual Risk Without Aggregate Uncertainty. *Journal of Economic Theory*, **35**, pp. 26-32.

[11] Gallego, G. and M. Hu. 2014. Dynamic Pricing of Perishable Assets under Competition. *Management Science*, **60**, pp. 1241-1259.

[12] Green, E.J. 1980. Non-cooperative Price Taking in Large Dynamic Markets. *Journal of Economic Theory*. **22**, pp. 155-182.

[13] Green, E.J. 1984. Continuum and Finite-player Noncooperative Models of Competition. *Econometrica*, **52**, pp. 975-993.

[14] Hopenhayn, H.A. 1992. Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, **60**, pp. 1127-1150.

[15] Housman, D. 1988. Infinite Player Noncooperative Games and the Continuity of the Nash Equilibrium Correspondence. *Mathematics of Operations Research*, **13**, pp. 488-496.

[16] Huang, M., R.P. Malhame, and P.E. Caines. 2006. Large Population Stochastic Dynamic Games: Closed-loop McKean-Vlasov Systems and the Nash Certainty Equivalence Principle. *Communications in Information and Systems*, **6**, pp. 221-252.

[17] Jouini, E., W. Schachermayer, and N. Touzi. 2006. Law Invariant Risk Measures Have the Fatou Property. *Advances in Mathematical Economics*, **9**, pp. 49-72.

[18] Jovanovic, B. and R.W. Rosenthal. 1988. Anonymous Sequential Games. *Journal of Mathematical Economics*, **17**, pp. 77-88.

[19] Judd, K.L. 1985. The Law of Large Numbers with a Continuum of I.I.D. Random Variables. *Journal of Economic Theory*, **35**, pp. 19-25.

[20] Kalai, E. 2004. Large Robust Games. *Econometrica*, **72**, pp. 1631-1665.

[21] Khan, M.A., K.P. Rath, and Y. Sun. 1997. On the Existence of Pure Strategy Equilibrium in Games with a Continuum of Players. *Journal of Economic Theory*, **76**, pp. 13-46.

[22] Levin, Y., J. McGill, and M. Nediak. 2009. Dynamic Pricing in the Presence of Strategic Consumers and Oligopolistic Competition. *Management Science*, **35**, pp. 32-46.

[23] Lin, K. and S.Y. Sibdari. 2008. Dynamic Price Competition with Discrete Customer Choices. *European Journal of Operational Research*, **197**, pp. 969-980.

[24] Mas-Colell, A. 1984. On a Theorem of Schmeidler. *Journal of Mathematical Economics*, **13**, pp. 201-206.

[25] Mertens, J.F. and T. Parthasarathy. 1987. Equalibria for Discounted Stochastic Games. CORE Discussion Paper No. 8750.

[26] Parthasarathy, K.R. 2005. *Probability Measures on Metric Spaces*. AMS Chelsea Publishing, Providence, Rhode Island.

[27] Perakis, G. and A. Sood. 2006. Competitive Multi-period Pricing for Perishable Products: A Robust Optimization Approach. *Mathematical Programming*, **107**, pp. 295-335.

[28] Reny, P.J. and M. Perry. 2006. Towards the Foundation for Rational Expectations Equilibrium. *Econometrica*, **74**, pp. 1231-1269.

[29] Sabourian, H. 1990. Anonymous Repeated Games with a Large Number of Players and Random Outcomes. *Journal of Economic Theory*, **51**, pp. 92-110.

[30] Schmeidler, D. 1973. Equilibrium Points of Nonatomic Games. *Journal of Statistical Physics*, **7**, pp. 295-300.

[31] Shapley, L.S. 1953. Stochastic Games. *Proceedings of the National Academy of Sciences*, **39**, pp. 1095-1100.

[32] Solan, E. 1998. Discounted Stochastic Games. *Mathematics of Operations Research*, **23**, pp. 1010-1021.

[33] Weintraub, G.Y., C.L. Benkard, P. Jeziorski, and B. van Roy. 2008. Nonstationary Oblivious Equilibrium. Working Paper, Columbia University.

[34] Weintraub, G.Y., C. L. Benkard, and B. van Roy. 2008. Markov Perfect Industry Dynamics with Many Firms. *Econometrica*, **76**, pp. 1375-1411.

[35] Weintraub, G.Y., C. L. Benkard, and B. van Roy. 2011. Industry Dynamics: Foundations for Models with an Infinite Number of Firms. *Journal of Economic Theory*, **146**, pp. 504-527.

[36] Xu, X. and W.J. Hopp. 2006. A Monopolistic and Oligopolistic Stochastic Flow Revenue Management Model. *Operations Research*, **54**, pp. 1098-1109.

[37] Yang, J. 2011. Asymptotic Interpretations for Equilibria of Nonatomic Games. *Journal of Mathematical Economics*, **47**, pp. 491-499.

[38] Yang, J. 2016. A Link between Semi-anonymous Sequential Games and their Large Finite Counterparts. *International Journal of Game Theory*, forthcoming.

[39] Yang, J. and Y. Xia. 2013. A Nonatomic-game Approach to Dynamic Pricing under Competition. *Production and Operations Management*, **22**, pp. 88-103.